Published on May 31, 2024

AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap

by Q. Vera Liao and Jennifer Wortman Vaughan

Abstract

The rise of powerful large language models (LLMs) brings about tremendous opportunities for innovation but also looming risks for individuals and society at large. We have reached a pivotal moment for ensuring that LLMs and LLM-infused applications are developed and deployed responsibly. It is paramount to pursue new approaches to provide transparency—a central pillar of responsible artificial intelligence (AI)—for LLMs, and years of research at the intersection of AI and human–computer interaction (HCI) highlight that we must do so with a human-centered perspective: Transparency is fundamentally about supporting appropriate human understanding, and this understanding is sought by different stakeholders with different goals in different contexts. In this new era of LLMs, we must develop and design approaches to transparency by considering the needs of stakeholders in the emerging LLM ecosystem, the novel types of LLM-infused applications being built, and the new usage patterns and challenges around LLMs, all while building on lessons learned about how people process, interact with, and make use of information. We reflect on the unique challenges that arise in providing transparency for LLMs, along with lessons learned from HCI and responsible AI research that has taken a human-centered perspective on AI transparency. We then lay out four common approaches that the community has taken to achieve transparency—model reporting, publishing evaluation results, providing explanations, and communicating uncertainty—and call out open questions around how these approaches may or may not be applied to LLMs. We hope this provides a starting point for discussion and a useful roadmap for future research.

Keywords: LLMs, generative AI, transparency, explainability, human-centered AI

Media Summary

While capturing intense public and academic enthusiasm, large language models (LLMs) suffer from a lack of transparency. We still do not have clear answers to even basic questions: What can an LLM do? How well can it do it? How exactly does it work internally? This poses challenges for different stakeholders. For instance, developers of LLMs and LLM-infused applications have difficulties debugging and determining responsible use of these models; end-users often lack a sufficient understanding of these models to interact effectively and know how much to trust them; and policymakers require more transparency to have effective oversight of LLMs. To facilitate research and practice on transparency of LLMs, in this article, we summarize lessons learned from artificial intelligence (AI) and human–computer interaction (HCI) research on AI transparency to encourage the community to embrace diverse approaches to transparency and center the development and evaluation of these approaches on people. We also call out unique challenges to achieving transparency for LLMs and discuss open questions about how existing transparency approaches may or may not be applied to LLMs.

1. Introduction

Hugely powerful large language models (LLMs) like GPT-4, LaMDA, and LLaMA are now being deployed in applications from search engines to code generation tools to productivity suites. These generative models are widely expected to have impact across industries, to change the way we engage in tasks like writing, programming, and design, and to potentially even reshape occupations in medicine, law, marketing, education, and beyond (Agrawal et al., 2022; Davenport & Mittal, 2022; DePillis & Lohr, 2023; Felten et al., 2023; P. Lee et al., 2023). As the chair of the U.S. Federal Trade Commission put it in a recent op-ed, "the full extent of generative AI's potential is still up for debate, but there's little doubt it will be highly disruptive" (L. M. Khan, 2023).

While the capabilities of LLMs are impressive, they also raise new risks (Bender et al., 2021; Kumar et al., 2023; OpenAI, 2023; Weidinger et al., 2022). Language models are found to encode biases (Abid et al., 2021; Rae et al., 2021), which risks propagating harmful discrimination, stereotypes, and exclusion at scale. They are widely known to 'hallucinate' information (Ji et al., 2022; S. Lin et al., 2022b; Maynez et al., 2020), producing outputs that are plausible—even convincing—but incorrect. They may project confidence about these hallucinated outputs, potentially contributing to automation bias, overreliance, or automation-induced complacency (Parasuraman & Manzey, 2010; Wickens et al., 2015). LLMs can generate harmful, sometimes toxic content, including hate speech and offensive language, or reveal sensitive information that threatens privacy or security. They can contribute—both intentionally and unintentionally—to the spread of misinformation (Buchanan et al., 2021; Kreps et al., 2022; Zhou et al., 2023). And in the longer term, LLMs may lead to environmental harms as well as socioeconomic harms, including the displacement and deskilling of workers across industries.

Given the anticipated impact that LLMs will have on both our day-to-day lives and society at large, it is critical that LLMs and LLM-infused applications be developed and deployed responsibly. One central component of responsible AI (artificial intelligence) development and deployment is *transparency*: enabling relevant stakeholders to form an appropriate understanding of a model or system's capabilities and limitations, how it works, and how to use or control its outputs. Developers of LLMs cannot debug their models, responsibly assess whether they are ready to launch, and enforce responsible and safe usage policies for their models without some understanding of their behavior and performance on different tasks. Business decision makers, designers, and developers building LLM-infused applications must be able to understand the LLM's capabilities and limitations in order to ideate and make decisions about whether, where, and how to use the model—potentially including how to fine-tune, prompt, or otherwise adapt the model to better fit their use case. End-users must be able to form a sufficiently accurate understanding of LLM-infused applications to control the application's behavior and achieve appropriate levels of trust and reliance. People impacted by LLMs or LLM-infused applications may require transparency in order to understand their options for recourse. Additionally, given the speed at which powerful new LLMs and their applications are being released and the growing concerns over potential harms, we should expect to see an increased demand for transparency around their development and inner workings from policymakers and third-party auditors aiming to regulate and oversee their use.

In recent years, we have witnessed the creation of a whole research field at the intersection of AI and human–computer interaction (HCI) that is focused on developing and evaluating different

approaches to achieve transparency. These approaches range from frameworks for documenting models and the data sets they are trained on (e.g., Arnold et al., 2019; Bender & Friedman, 2018; Crisan et al., 2022; Gebru et al., 2021; Holland et al., 2018; Mitchell et al., 2019) to techniques for producing explanations of individual model outputs (e.g., Koh & Liang, 2017; Lundberg & Lee, 2017; Ribeiro et al., 2016; Russell, 2019; Ustun et al., 2019) to approaches for communicating uncertainty (e.g., Bhatt et al., 2021; Dhami & Mandel, 2022; D. Wang et al., 2021) and beyond. There is no one-size-fits-all solution. In the case of LLMs, the needs of an application developer engaging in ideation are probably different from those of a writer who is using an LLM-infused application to edit a novel or a public figure who is concerned about how their life is presented by an LLM-infused search engine. In our own work (Liao & Varshney, 2021; Vaughan & Wallach, 2021), we have argued for the importance of taking a human-centered perspective on transparency—designing and evaluating transparency approaches with stakeholders and their goals in mind. We believe that this is even more important in the era of LLMs, when the diversity of stakeholders and their experience levels, contexts, goals, and transparency needs is greater than ever.

In this article, we map out a human-centered research roadmap for transparency in this new era. We first reflect on the unique challenges that arise in providing transparency for LLMs compared with smaller-scale, more specialized models that have traditionally been the focus of AI transparency research. We reflect on lessons learned from HCI and Responsible AI/FATE (fairness, accountability, transparency, and ethics) research that is concerned with human needs of, interactions with, and impact from AI transparency. We then lay out common approaches, including techniques and artifacts, that the community has taken to achieve transparency and call out open questions around how they may or may not be applied to LLMs.

We note that there is no agreed-upon definition of transparency, and indeed, transparency has been recognized as a multifaceted concept. In this article, we adopt a focus on *informational* transparency—essentially, what information about a model (or system building on that model) should be disclosed to enable appropriate understanding—which has been emphasized within the machine learning (ML) research community and in industry practice, though we note that there are other perspectives, such as the normative and relational dimensions of transparency, that have been studied in the broader literature (Felzmann et al., 2020; Meijer, 2013). Some of the approaches we cover, such as model reporting, are primarily aimed at supporting a *functional* understanding of *what* the model (or system) can do, often by exposing the goals, functions, overall capabilities, and limitations. Others, like the explanations frequently explored in the explainable AI (XAI) and interpretable ML communities, are primarily aimed at supporting a *mechanistic* understanding of *how* the model (or system) works, by disclosing the parts and processes (Lombrozo, 2012). We believe that both understandings play important roles and the appropriate form of transparency will depend on the stakeholder and the goal that they wish to achieve.

Finally, we note that many of the challenges, lessons learned, potential approaches, and open problems that we call out in this article apply not only to LLMs but to other large-scale generative models, including multimodal models that allow for both textual and visual input or output. While we adopt the narrower focus on LLMs for simplicity, we encourage additional research on transparency for these other models.

2. What Makes Transparency for LLMs Challenging?

To ground the discussion in the remainder of the article, we first explore the unique characteristics of LLMs and the emerging patterns of their usage that are likely to make it more challenging to achieve transparency compared with the smaller-scale, specialized models that AI transparency research has traditionally dealt with. We start by providing some brief background on LLMs and establishing some terminology that we will use in the rest of the article.

2.1. Background on LLMs

An LLM, like any language model, predicts the conditional probability of a token—which might be a character, word, or other string—given its preceding context and, in the case of bidirectional models, its surrounding context (Bengio et al., 2003; Radford et al., 2019). Present-day LLMs are based on modern neural network self-attention architectures like the transformer (Vaswani et al., 2017) with hundreds of billions or even more than a trillion parameters (Ganguli et al., 2022). While earlier models were trained on data sets of moderate size, LLMs are trained on data sets of massive scale, with hundreds of billions or even more than a trillion tokens (Borgeaud et al., 2021; Hoffmann et al., 2022), requiring many orders of magnitude more compute time. This makes LLMs vastly more sophisticated and expressive than their predecessors.

While a basic pretrained LLM model can be viewed as a 'general purpose' next-word predictor, LLMs can be adapted to exhibit or suppress specific behaviors or to perform better on specific tasks like text summarization, question answering, or code generation. One common approach is fine-tuning, in which the model's parameters are updated based on additional, specialized data (e.g., Devlin et al., 2019; Howard & Ruder, 2018; C. Lee et al., 2020; Radford et al., 2018). A popular technique for fine-tuning is reinforcement learning from human feedback (RLHF) in which human preferences are used as a reward signal (Christiano et al., 2017; Ouyang et al., 2022). Another approach is prompting or prompt engineering, in which natural language prompts—often containing examples of tasks (for few-shot prompting/in-context learning) or demonstrations of reasoning (for chain-of-thought prompting)—are provided to the model to alter its behavior without making any changes to the model's internal parameters (e.g., Brown et al., 2020; J. Liu et al., 2022; Shin et al., 2020; Wei, Wang, et al., 2022). The adapted model can be incorporated into applications such as chatbots, search engines, or productivity tools. Models can also be augmented with the ability to call on external models, tools, or plugins (Mialon et al., 2023), for example, querying an information retrieval system to ground their output or controlling and receiving feedback from a physical robot.

It is important to note that the party adapting the model or building the application is frequently not the same party who built the underlying pretrained LLM, and may only be able to access the LLM through an application programming interface (API). A model may also be adapted more than once by different parties; for instance, a base model may be fine-tuned using RLHF by its creators, fine-tuned on domain-specific data by application developers, and then adapted via in-context learning by end-users. When we talk about transparency, we must keep in mind whether we are referring to transparency about the pretrained LLM, an adapted LLM, or the application using the pretrained or adapted model (LLM-infused application). We aim to call out which of these we are referring to when it is not clear from the context.

2.2. Challenges for Achieving Transparency

There are several characteristics of LLMs and their usage that pose challenges for transparency. The list we lay out here is not meant to be exhaustive, but to provide context for later discussion.

Complex and Uncertain Model Capabilities and Behaviors. LLMs can perform an astonishingly wide variety of tasks in different contexts (Bommasani et al., 2021). Unlike classical machine learning models where there is typically a well-defined structure of inputs and outputs, LLMs are more flexible. The capabilities of LLMs—sometimes also referred to as use cases (Ouyang et al., 2022) or tasks (Liang et al., 2022) in the literature—include question answering, dialogue generation, sentence completion, summarization, paraphrasing, elaboration, rewriting, classification, and more. Researchers are now additionally identifying 'emergent capabilities' of LLMs—like performing arithmetic or chain-of-thought reasoning—that are not present in smaller-scale models but emerge at scale (Wei, Tay, et al., 2022). Furthermore, as described above, the precise behavior and capabilities of an LLM can be steered through approaches like fine-tuning and prompting. All of this contributes to "capability unpredictability" (Ganguli, Hernandez, et al., 2022), the idea that an LLM's capabilities cannot be fully anticipated, even by the model's creators, until its behavior on certain input is observed.

Additionally, present-day LLMs exhibit unreliable behaviors. Their responses can change with updates, the details of which are often not made transparent by LLM providers. Depending on the sampling strategy used (Holtzman et al., 2020), outputs can be nondeterministic in the sense that the same prompt leads to a different response when input to the model again. They can misinterpret a prompt in unpredictable ways and respond inconsistently to a type of prompt, making the behavior of adapted models difficult to predict. These unreliable behaviors can make it challenging, if not impossible, to gain a generalized understanding of the model's behavior.

Massive and Opaque Architectures. Given the complexity and massive scale of present-day LLMs, there are currently no techniques that would provide us with a complete picture of the knowledge reflected in a model or the reasoning that is used to produce its output (Bowman, 2023). The mechanism of the transformer architecture underpinning LLMs is yet to be fully understood, even among experts, and some techniques that initially appear promising for interpreting the behavior of LLMs, such as looking at attention weights or perturbing inputs, can be misleading (Bolukbasi et al., 2021; Feng et al., 2018; Jain & Wallace, 2019). A more unique challenge with LLMs is the massive scale of the training data and diverse sources from which it is pulled—for example, Common Crawl and Wikipedia, often with no specific topics or formats targeted, and no thorough documentation about how the data set was developed (M. Khan & Hanna, 2022). This makes it challenging, if not impossible, to understand what went into an LLM's training. There are currently no established answers, even in the research community, to questions such as precisely how and why these models work as well as they do, why they can or cannot perform certain tasks, and how characteristics of the training data impact model capabilities. Furthermore, the sheer size can make it challenging to develop and operationalize transparency approaches (e.g., due to constraints on method scalability or computing resources).

Proprietary Technology. An elephant in the room that will inevitably inhibit attempts at transparency for LLMs is the proprietary nature of the models. Currently, while the efforts for developing open source LLMs are growing (BigScience Workshop, 2022; Touvron, Lavril, et al., 2023; Touvron, Martin, et al., 2023), most of the powerful LLMs were developed at large technology companies or other nonacademic organizations. They are either released through APIs or completely proprietary, making it impossible to access their inner workings (e.g., weights and parameters). In many cases, details such as the size, makeup and provenance of the training data, the number of parameters, and the resources required to train the model are also not shared publicly. In essence, then, such models can only be probed in a black-box manner, which may not be

sufficient to meet the transparency requirements for stakeholders, and poses challenges for the research community to develop transparency approaches. Addressing this fundamental challenge may not be possible without regulatory efforts that enforce transparency requirements on LLM creators and providers.

New and Complex Applications. End-users may not interact with LLMs directly, but rather through LLM-infused applications. Emerging applications include general and specialized chatbots, web search, programming assistants, productivity tools such as for writing support or presentation generation, and text analysis tools such as for customer insights discovery. As LLMs' capabilities continue to be discovered, we can only expect the number and variety of LLM-infused applications to grow. While any opacity of the model will likely trickle down to hinder the transparency of the applications built on it, as mentioned above, the transparency requirements for LLM-infused applications will be different from the requirements for the model as they serve a different set of stakeholders. Furthermore, just as the models themselves are flexible, the use cases for LLM-infused applications can be flexible and open-ended. For example, an LLM-infused search engine may be used to plan a trip, research a report, or write a poem—use cases that reflect different needs in terms of accuracy, verifiability of output, and likely the required transparency approaches.

To further complicate transparency around LLM-infused applications, such applications may not be built on a single LLM, but may involve many interacting models and tools. For example, auxiliary LLMs can be used to augment the output or expand the capabilities of a primary LLM. LLMs can be embedded in a complex system to operate other models or external services, for example, through plugins, allowing them to perform tasks like ordering groceries or booking flights with no human in the loop. An application may also include other components like input or output filters. For example, an LLM-infused search engine may rely on results obtained from a traditional search engine to 'ground' its responses. Changes to any component can change the behavior of the application, making it more difficult to understand its behavior. Approaches to transparency must therefore take into account all components and how they fit together rather than focusing on an LLM in isolation.

Expanded and Diverse Stakeholders. As the number of LLM-infused applications grows and popular applications such as LLM-infused search engines expand their user bases, a larger number of people—diverse along many dimensions—will interact with or be impacted by LLMs. Research in AI transparency typically considers stakeholder groups like data scientists and model developers, business decision makers, regulators and auditors, end-users, and impacted groups (i.e., the people who are directly or indirectly affected by a model or application) (Hong et al., 2020; Liao & Varshney, 2021; Vaughan & Wallach, 2021). The use of LLMs may introduce new stakeholder groups with unique transparency needs. For example, it is increasingly common for product teams to have dedicated prompt engineers—a role that, until recently, did not exist—to streamline tasks, evaluate models, or contribute to model adaptation. As another example, as LLMs are increasingly used for productivity support to augment people's writing, we must consider both the creators of LLM-assisted articles and the consumers of these articles as 'users' of the LLM's outputs and support both groups' transparency needs. Meanwhile, we must support any subjects referred to in the articles as 'impacted groups.'

As the pretrained nature of LLMs lowers the barrier to using and building on AI capabilities, we believe application developers—including those working on model adaptation—will become a significant group and diverse in itself, potentially including developers, entrepreneurs, product

managers, designers, or essentially anyone. In some cases, the line between application developers and end-users may be blurred. Consider, for example, a writer who experiments with using an LLM for writing support. This writer might benefit from model transparency to assess the LLM's suitability for different writing tasks and identify effective ways to adapt the model for each task.

Recent research has begun to inquire about the ecosystem of LLMs and the roles in it (Bommasani et al., 2021), from data creation, curation, model training, and model adaptation through to deployment. Identifying these LLM stakeholder roles and supporting their role-, task- and context-specific transparency needs will be of primary importance for AI transparency research.

Rapidly Evolving and Often Flawed Public Perception. Effective approaches to transparency should take into account the receivers' existing perception of what the model or system can or cannot do, and how it works—often referred to as their mental model (Gentner & Stevens, 1983; Johnson-Laird, 1983; Norman, 1987). This is especially challenging for LLMs as their public perception is still evolving and shaped by complex mechanisms including mass media, marketing campaigns, ongoing events, and design choices of popular LLM-infused applications. The natural language modality also contributes to a unique set of challenges: people may be more likely to assign humanlike attributes to the model and have corresponding expectations (Nass & Moon, 2000), and even subtle language and communication cues can have profound impact on people's mental model (Abercrombie et al., 2023). Recent studies show that people already have flawed mental models about LLMs, such as incorrect perceptions of how their output differs from human-written texts (Jakesch et al., 2023). Interacting with LLMs with a flawed mental model can lead to misuse, unsafe use, over- and underreliance, deception, privacy and security threats, and other interaction-based harms (Weidinger et al., 2022). Flawed public perceptions can be attributed to a lack of accurate, comprehensive, and responsible information. In addition to incorporating transparency approaches, the organizations creating LLMs and LLM-infused applications and the research community more broadly should reflect on the implications of the way they communicate with the public. For example, the use of ill-defined, catch-all phrases such as 'general purpose model' or inappropriate anthropomorphizing may hinder accurate public perception of LLMs.

Organizational Pressure to Move Fast and Deploy at Scale. Lastly, we note that there are organizational challenges that may hinder the development and adoption of transparency approaches beyond the proprietary nature of LLMs. Responsible AI efforts are often in tension with pressures to release products quickly and to scale up across geographies, use cases, and user bases (M. Madaio et al., 2022; M. A. Madaio et al., 2020; Rakova et al., 2021), a kind of "scale thinking" (Hanna & Park, 2020). Given the speed at which research and product breakthroughs are occurring and the vast financial stakes, companies are incentivized to move at a pace that is unusual to witness even in the technology industry to be first to market—what some media outlets have dubbed an "AI race" or "AI arms race" (Chow & Perrigo, 2023; Grant & Weise, 2023). The organizations building LLMs and LLM-infused applications will need to take extra steps to ensure that transparency and other responsible AI considerations are not lost in the process, which may require enhanced internal governance or external regulatory requirements in addition to organizational incentives for the individuals working in this space.

3. What Lessons Can We Learn From Prior Research?

In this section, we reflect on lessons from the HCI and Responsible AI/FATE research communities, which tend to take a human-centered perspective on transparency. While many technical transparency approaches (to be discussed in the next section) have been developed to deal

with what information about the model can be disclosed, the human-centered perspective focuses on how people *use* and cognitively *process* transparency information. Knowledge from this human-centered perspective should drive the development of transparency features, which concern not only the model-centered techniques but also the interfaces.

3.1. Transparency as a Means to Many Ends: A Goal-Oriented Perspective

Within the HCI community, researchers have attempted to guide the development and evaluation of transparency approaches by digging into the reasons why people seek information (M. Langer et al., 2021; Suresh et al., 2021). This goal-oriented perspective resonates with studies of human explanations from cognitive science, philosophy, and psychology (Lombrozo, 2016; Miller, 2019), where it is recognized that seeking explanations and achieving understanding is often a means to an end for downstream cognitive tasks like learning, decision-making, trust development, and diagnosis.

This goal-oriented perspective has led to works developing taxonomies of common goals that people seek explanations for (Liao et al., 2020, 2022; Suresh et al., 2021) and empirical studies to delineate common transparency goals of stakeholder groups such as data scientists (Bhatt et al., 2020; Hong et al., 2020) and designers (Liao et al., 2023). By focusing on the goals (as opposed to the low-level application or interface types), this perspective provides a useful level of abstraction to consider people's different transparency needs according to their usage of the information. For example, Suresh et al. (2021) lay out a set of common goals that people seek AI explanations for, including improving a model, ensuring regulatory compliance, taking actions based on model output, justifying actions influenced by the model, understanding data usage, learning about a domain, and contesting model decisions. By focusing on the normative goals that explanations should help people to achieve understanding for, the Information Commission Office of the United Kingdom's guidance on "explaining decisions made with AI" (Information Commissioner's Office, 2020) specifies six types of explanations: rationale, responsibility, data, fairness, safety and performance, and impact.

The goal-oriented perspective also has several practical implications for developing human-centered transparency approaches. First, whether a transparency approach is effective should be evaluated by whether it successfully facilitates a stakeholder's end goal. This means that not all situations require the same level of transparency (e.g., a low-stakes application such as generating poetry for fun may require little transparency). This also requires articulating end goals up front in order to choose criteria for evaluating transparency approaches. As an example, in our own work with collaborators, we focused on data scientists' goal of model debugging and evaluated two common techniques from the interpretable machine learning literature in terms of how well they help data scientists identify common problems with training data sets and the resulting models, finding evidence that the techniques may hamper the debugging goal by leading to overtrust or overconfidence about the model (Kaur et al., 2020). Second, achieving an end goal may require information beyond details of the model, such as information about the domain and the social-organizational context the model is situated in (Ehsan et al., 2021), and hence require holistic support with information tailored to the task at hand and integrated into the application interface.

What are the new transparency goals for LLMs? The new ecosystem and novel applications of LLMs call for investigations into what are the new types of common stakeholder goals that require transparency. For example, there may be heightened needs for supporting ideation, model

adaptation, prompting, and discovering risky model behaviors. New transparency approaches for LLMs should be developed and evaluated in terms of how well they help achieve these goals.

3.2. Transparency to Support Appropriate Levels of Trust

Although transparency has often been embraced within the tech industry as a mechanism to build trust, recent HCI research has taken the position that transparency should instead aim to help people gain an *appropriate* level of trust (Bansal et al., 2023)—enhancing trust when a model or application is trustworthy, and reducing trust when it is not. While relevant to many use cases of transparency, achieving an appropriate level of trust is especially critical for end-users to harness the benefits of AI systems without overrelying on flawed AI outputs.

Empirical studies on the relationship between transparency and user trust have painted a complex picture. In particular, a wave of HCI studies repeatedly showed that AI explanations can lead to overreliance—increasing people's tendency to mistakenly follow the AI outputs even when they are wrong (Bansal et al., 2021; V. Chen et al., 2023; Poursabzi-Sangdeh et al., 2021; X. Wang & Yin, 2021; Y. Zhang et al., 2020). Understanding this pitfall of AI explanations requires paying attention to people's cognitive processes. Researchers have attributed this difficulty to detect model errors from popular forms of AI explanations to their complexity and incompatibility with people's reasoning process (V. Chen et al., 2023), as well as to the heuristics and biases that people bring into their cognitive processes, such as an inclination to superficially associate an AI system being explainable with it being trustworthy (Liao & Sundar, 2022). Studies of other transparency approaches have also reported nuanced results (Rechkemmer & Yin, 2022; Schmidt et al., 2020; Yin et al., 2019; Y. Zhang et al., 2020). For example, while one study demonstrates that communicating uncertainty is more effective than providing explanations in supporting appropriate trust (Y. Zhang et al., 2020), another study suggests that people's trust level is more likely to be dominated by aggregate evaluation metrics such as accuracy (Rechkemmer & Yin, 2022).

Which approaches to transparency can best support appropriate trust of LLMs and how? There is a need to disentangle the relationship between trust and transparency for LLMs through both better conceptualization and careful empirical investigations. For the former, recent FATE literature has begun to unpack trust as a multifaceted and multiloci concept (Jacovi et al., 2021; Liao & Sundar, 2022). For LLMs, people's locus of trust can be at the base model, the LLM-infused application, the application provider (e.g., based on brand), or specific application functions or types of outputs, each of which may require different kinds of transparency support but also be intertwined with other loci. For example, people need to understand and reconcile that LLMs are powerful technologies but may not be used reliably for a certain application function. For empirical investigations, there is extensive literature on measuring trust on which to build (Vereschak et al., 2021), though it remains a challenge in practice (Bansal et al., 2023), and even more so with the complex, dynamic, and multiloci nature of trust around LLMs. Furthermore, evaluating appropriate trust requires further unpacking the actual 'trustworthiness' of a model or system and what counts as 'appropriate,' both of which remain open questions for LLMs.

3.3. Transparency and Control Often Go Hand-in-Hand

Many of the end goals we discussed in Section 3.1, such as improving or contesting the model and adapting data usage, can only be achieved by having both a good understanding of the model and appropriate control mechanisms through which to take action. Indeed, transparency and control have long been studied together in HCI as intertwined design goals for effective user experience

(M. K. Lee et al., 2019; Wu et al., 2022). This is well reflected in the interdisciplinary area of interactive machine learning (iML) (Amershi et al., 2014)—learning interactively through feedback from end-users—and related areas such as machine teaching (Carney et al., 2020; Simard et al., 2017). These paradigms simultaneously ask what information about a model should be presented to users and what forms of input or feedback users should be able to give in order to steer the model. We believe current work on training, adapting, and building applications around LLMs can take valuable lessons from these lines of research. More recent HCI studies on algorithmic transparency also highlight that providing transparency without supporting control leaves users frustrated, while effective, efficient, and satisfying control cannot be achieved without transparency (Smith-Renner et al., 2020; Storms et al., 2022). More critically, scholars have called out the risk of algorithmic transparency without paths for actionability and contestibility as creating a false sense of responsibility and user agency (Ananny & Crawford, 2018; Kluttz et al., 2020; Lyons et al., 2021).

How can different approaches to transparency contribute to better control mechanisms for LLMs? While safety and control have become central topics in research and practices around LLMs (Keskar et al., 2019; Li et al., 2023), the role of transparency is less emphasized. We encourage the community to consider the role of transparency in establishing better mechanisms for control and enabling more participatory and inclusive approaches that allow stakeholders to understand and then steer LLM behavior.

3.4. The Importance of Mental Models

People's existing understanding of a system impacts what information they seek for transparency and how they process the information. This is often studied in HCI work through the concept of a mental model—one's internal representation of a system based on their experiences, whether direct or indirect, with the system. A good mental model should be both accurate and complete, as it is the foundation for effective, efficient, and satisfying interactions with a system (Norman, 2014). HCI research also differentiates between a functional (shallow) mental model—knowing what a system can be used for and how to use it—and a structural (deep) mental model—knowing how and why the system works (Kulesza et al., 2012). Transparency approaches for functional and mechanistic understandings can be seen as supporting these two aspects of mental models, respectively. However, since mental models are shaped by continuous interactions with a system, some researchers have argued that notions like the 'interpretability' of an AI system need to be considered as evolving through dynamic and situated system interactions rather than considered in the context of a single intervention like the introduction of documentation or explanations (Thieme et al., 2020).

We highlight several ways that transparency approaches should consider people's mental models. First, transparency approaches should be designed to support different stakeholders in building a good mental model. It may therefore be appropriate for evaluations of transparency approaches to incorporate assessments of mental model accuracy and completeness, for example, by analyzing people's comments or answers to questions about their beliefs about a model or application's function and structure (Eslami et al., 2016; Gero et al., 2020; Grill & Andalibi, 2022; Kulesza et al., 2012). Second, transparency approaches should account for people's existing mental models, and focus on closing the necessary gaps to allow them to achieve their end goal (Eiband et al., 2018). This means that approaches to transparency should avoid conveying redundant information that people already have in their mental models, but more importantly, aim to correct flawed mental models. However, it is known that a mental model, once built, is often difficult to shift even if

people are aware of contradictory evidence (Wilfong, 2006), which may present a significant challenge for transparency approaches to be effective. This highlights the importance of responsible communication (e.g., in marketing material and media coverage) to accurately shape the public perception around new technologies like LLMs. In addition, Norman (2014) noted that people's mental models are often incomplete, unstable, have unclear boundaries (e.g., mixing up different parts of the systems), and favor simple rules, all of which may pose challenges for transparency approaches to help people build an appropriate understanding.

How can we unpack people's mental models of LLMs and support forming better mental models? Just as it is difficult to characterize the capabilities and limitations of LLMs given their scope and capability unpredictability, it is difficult to characterize people's mental models of them. More research is also needed to understand the general mental models that people already have of LLMs, especially in response to their unique characteristics such as humanlike language capabilities and unreliable behaviors (e.g., hallucinating and nondeterministic output). Moreover, HCI research has traditionally dealt with mental models at the system level, while people's mental models of an LLM-infused application could be muddled by the blurred boundaries between the pretrained model, the adapted model(s) used in the application, and the application itself. While it remains critical for transparency approaches to aim to correct flawed mental models and build accurate and complete mental models, the field may need foundational work on how to characterize, assess, and offer opportunities to build and shift mental models of LLMs.

3.5. How Information Is Communicated Matters

HCI research on AI transparency is often concerned with not only what information to communicate about a model, but how to communicate it. Work has explored ways of communicating performance metrics(Görtler et al., 2022), explanations of model outputs (Hadash et al., 2022; Lai et al., 2023; Szymanski et al., 2021), and uncertainty estimates (Kay et al., 2016), as well as how to frame the model's output itself in order to appropriately shape people's mental model (e.g., whether to use certain terms like 'risk' [Green & Chen, 2021]). Such information can be communicated through different modalities (e.g., by a visualization or in natural language), at different levels of precision or abstraction, framed using different language, supplemented with different information to close any gaps in understanding, and through various other visual and interactive interface designs. These choices of communication design can significantly impact how people perceive, interpret, and act on the information provided.

An effective design should be guided by the ways that people process information cognitively and socially. For example, a line of HCI research explored more user-friendly visualization designs to overcome the trouble that people often have understanding statistical uncertainty and the cognitive biases they bring (Fernandes et al., 2018; Kay et al., 2016). In light of the difficulty of reasoning about the complex explanations produced by some AI explainability techniques, HCI research has explored how to present explanations in more human-compatible ways (Hadash et al., 2022; Szymanski et al., 2021). In our recent work with collaborators (Lai et al., 2023), we argue that people engage in two processes to produce explanations (Malle, 2006): an information-gathering process in which they come up with a set of reasons, and a communication process to present reasons, often selectively tailored to the recipient. Explainability techniques that focus on revealing the inner workings of a model are typically only concerned with the former. We then propose a framework to tailor these explanations by learning the recipient's preferences as a selective

communication strategy, and empirically demonstrate that these selected explanations are easier to process and better at helping people detect model errors in an AI-assisted decision-making task.

What are the new opportunities and challenges for communicating information during interactions with LLMs? The natural language modality of LLMs has significant implications for the communication aspect of transparency. For example, instead of presenting a numerical score for uncertainty, LLM-infused applications like chatbots can express uncertainty by using hedging language or refusing to answer a question. This behavior can now be built into the adapted model directly through fine-tuning or prompting (S. Lin et al., 2022a), making it potentially harder to precisely control and interpret the communication. Meanwhile, as decades of HCI research on chatbots and conversational interfaces suggest, people's perceived utility of these technologies can be shaped by a wide range of communication, social, and linguistic behaviors such as how the agents introduce and clarify their capabilities, take initiatives, repair errors, and respond to chit-chat requests, and even their language style (e.g., Ashktorab et al., 2019; Avula et al., 2022; Langevin et al., 2021). We believe more research is needed to distill principles to effectively communicate necessary information about the model's capabilities, limitations, and mechanisms during natural language interactions, as well as to establish reliable approaches for LLMs to follow these principles.

3.6. Limits of Transparency

Last but not least, we call attention to some critiques on the limits of transparency offered by FATE and science and technology studies (STS) scholars (Ananny & Crawford, 2018; Knowles, 2022; Lima et al., 2022). First, related to several arguments throughout the article, model-centric transparency without ensuring human understanding or meaningful effects on people's end-goals—"seeing without knowing" (Ananny & Crawford, 2018)—loses its purpose, and worse, can create a false sense of power and agency. Second, transparency can be misused to shift accountability and place burdens on users, and can even be used to intentionally occlude information. Those users without the necessary technical background and training to make sense of the provided information may face higher burdens. This is a warning to the field to pay attention to the consumability of transparency approaches and to seek alternative paths to ensure accountability. Lastly, transparency approaches can lead to harms if used maliciously or inappropriately, including exploiting user trust and reliance. In addition, transparency may present tensions with other stakeholder goals such as security and privacy, and such tensions may create disparities of benefit and cost between stakeholder groups (e.g., harming the data subject while benefiting the developer).

When is transparency not enough, and what else do we need? More research is needed to understand the limits of transparency for LLMs and how to properly hold the organizations building and deploying LLMs and LLM-fused applications accountable. The latter may require policy and regulatory changes, in addition to new approaches for external auditing (Mökander et al., 2023).

4. What Existing Approaches Can We Draw On?

The ML and HCI research communities have explored a variety of approaches to achieving transparency, including model and data reporting, publishing the results of evaluations, generating explanations, and communicating uncertainty. In this section, we briefly review these approaches and explore the extent to which they may or may not be applicable in the context of LLMs, while calling out needs specific to stakeholders of LLMs and open questions that arise. We note that this is not an exhaustive list of approaches and encourage the research community to explore new

categories of approaches that can help people achieve functional and mechanistic understandings of LLMs. For example, along the way, we suggest areas such as tools to support model interrogation and communicating output-specific risk and safety concerns.

4.1. Model Reporting

Documentation has become a building block for responsible AI in industry practice. Standardized documentation frameworks have been proposed to encourage both reflection and transparency around models (Mitchell et al., 2019), AI services (Arnold et al., 2019), and training and evaluation data sets (Bender & Friedman, 2018; Gebru et al., 2021; Holland et al., 2018). For example, the model cards framework (Mitchell et al., 2019), a popular framework for model reporting that has been adopted by companies like Google and Hugging Face, specifies comprehensive information that should be reported about a model, including a description of its inputs and outputs, the algorithm used to train it, the training data, additional development background, the model's intended uses, and ethical considerations. The framework emphasizes the inclusion of quantitative model evaluation results (more on that in the next section), including disaggregated evaluations (Barocas et al., 2021), in which results are broken down by individual, cultural, demographic, or phenotypic groups, domain-relevant conditions, and intersections of multiple groups or conditions. Disaggregated evaluation can help identify fairness issues, and also assist stakeholders in identifying when or where the model is suitable or reliable to use. In short, good documentation can help stakeholders who are building on a model or data set assess its suitability for their purpose and avoid misuse. It can also provide the necessary context for end-users, impacted groups, regulators, and auditors to understand how models and systems are being built and deployed.

While celebrated as an approach to providing transparency, creating good documentation remains challenging in practice. In our prior work with collaborators, we found that practitioners tasked with documenting a data set they worked with struggled to make the connection between the information that they were asked to include and its implications for responsible AI, were unsure of the appropriate level of detail to include and who the target audience was, and in some cases were uncertain about what even counts as a data set (Heger et al., 2022). Some stakeholders also struggle to consume existing forms of documentation. For example, designers or analysts without formal training in machine learning can find standard documentation to be too technical, and the lengthy textual format to be cumbersome (Crisan et al., 2022; Liao & Sundar, 2022).

What information is needed to characterize the functional behavior of an LLM? In principle, existing model reporting frameworks could be applied as is to LLMs. However, some of the information categories in a standard model card would be difficult to pin down due to the 'general purpose' positioning of LLMs and the uncertainty surrounding their capabilities. Even providing basic details such as what the input and output spaces of an LLM or LLM-infused application are, and the mapping between inputs and outputs, can be an elusive task. Currently, it is common for LLM providers to instead provide a description of intended use cases (like 'summarization' or 'creative and collaborative writing') or demonstrations of example prompts and responses. While this information can be a useful component of model reporting, it can also be misleading or, in some cases, even deceptive, since cherry-picked examples can shape user and public perception in a skewed way. This raises questions about how these examples should be selected and who should select them.

While we elaborate on performance reporting in the next section, we call out two other important categories in the model cards framework that are currently missing or incomplete for most LLMs:

training data and development background. Besides the incentive for organizations to keep this information proprietary, we must recognize that there are open questions about how to provide such information given the complexity of LLMs and unique aspects of their training processes.

For data transparency, as discussed in Section 2, the data sets used to pretrain base models are unprecedentedly massive in scale and pulled from diverse sources. Conveying their full scope and makeup is impossible, but there may be ways of distilling the most critical characteristics of these data sets to provide a basic understanding of what goes into the models. Different issues arise when considering the data sets used for model adaptation. For example, as companies engage in user data collection for the purpose of fine-tuning models, they must pay due diligence to the transparency of their user data handling, including privacy.

For development background, besides standard information such as the choice of algorithms, architecture, and parameters, LLM providers should include additional details on the training process. For example, an emerging practice is for LLM development to include some sort of 'alignment' effort to make the model more usable or safe (e.g., producing less toxic or harmful content). This can be done using human feedback through RLHF (Ouyang et al., 2022) or by having the model critique itself based on human-specified rules or principles (Bai et al., 2022). Given that LLMs' behaviors can be governed by these alignment efforts, it is especially important to make them transparent to allow the public and regulatory bodies to understand, scrutinize, and iterate on them.

What do different (and new) types of stakeholders need from model reporting frameworks? In light of the lessons discussed in Section 3, we recommend more research on the fundamental question of what different stakeholders want to know—and what they *should* know—about the model, along with a careful examination of how different forms of information shape their perception and usage of LLMs. As LLMs change the ML product development and deployment lifecycles, we may need to revisit the positioning of model reporting and consider new types of frameworks that address the specific needs of new stakeholder groups. For example, as discussed above, the LLM ecosystem introduces a new stage of model adaptation through fine-tuning, prompting, or other techniques. This adaptation may be performed by the original model builder, an application developer, or in some cases, directly by end-users. To date, there has been little or no research on these stakeholders' transparency needs when adapting the model, or about how they should transparently convey information about model adaptations to other parties.

What is needed beyond static documentation? Lastly, we call out that model reporting should not be limited to static, textual documentation or a basic 'card' format. Any formats or features that provide functional information about the model and shape people's understanding can contribute to model reporting. These may include FAQ pages, landing or onboarding pages, or even media communication describing the model. All such features can benefit from standardization and, where appropriate, regulation.

Following recent HCI and FATE studies investigating how to design effective documentation interfaces (Crisan et al., 2022; Liao et al., 2023), we suggest that those designing model reports for LLMs should explore more interactive features. For example, prior works have explored interfaces for uploading, customizing, and slicing input data to generate customized reports and visualize input-output spaces. Interactive interfaces are particularly suitable for LLMs for several reasons. First, interactive features can better support information navigation and consumption to accommodate LLM stakeholders from diverse backgrounds. Second, interaction allows for

experienced affordance and interrogation to understand LLMs' complex capabilities and behaviors that could be difficult to capture with textual descriptions. Lastly, as our study with collaborators on designers' use of model documentation suggests (Liao et al., 2023), static documentation presents significant gaps for contextualizing the model capabilities and limitations for one's own setting. It will be impossible for documentation creators to anticipate every downstream use case of LLMs. Instead, stakeholders should be provided with opportunities to interrogate the model with their own input data, capabilities of interest, hypotheses, and questions.

4.2. Publishing Evaluation Results

While evaluation results are often included as one component of a model report, we believe that publishing evaluation results is an important and complex enough topic that it deserves a separate discussion. Beyond model reports, evaluation results may also be published by third-party auditors or researchers for the purpose of ensuring compliance with regulations or standards, benchmarking, or exposing model limitations or potential harms. As discussed, evaluation can happen at the aggregate level or can be disaggregated by groups or conditions of interest. Evaluations may also be performed on a model or on the full system into which it is incorporated. While performance quality (e.g., some notion of accuracy) is often the primary focus of an evaluation, evaluations may also consider fairness (through disaggregated evaluations or using specific fairness metrics), robustness, efficiency, or other characteristics of a model or system's behavior, including how they impact end-users.

We note that the ML and natural language processing (NLP) communities have long dealt with the challenges of evaluating the performance of generative models (Sai et al., 2022). Until recently, natural language generation (NLG) evaluations have focused on tasks that specialized NLG models commonly perform, such as machine translation, abstractive summarization, question answering, and dialogue generation. For tasks that involve classification, standard performance metrics relying on exact matching with ground-truth labels like accuracy, precision, and recall can be used. In contrast, when the output space is open-ended and complex, as it often is for generative models, it becomes necessary to rely on more sophisticated performance metrics for word-based or embedding-based matching—for example, ROUGE score (C.-Y. Lin, 2004) or BERTscore (T. Zhang et al., 2019)—and more complicated (but often flawed) ways to obtain a 'ground-truth' reference to compare against. In practice, ground-truth data are often chosen either because they are conveniently available, such as using the 'highlights' of news articles as the ground-truth for summarization (See et al., 2017), or generated by crowd workers. Recently there has been a wave of data auditing work questioning the assumptions behind and quality of some widely used evaluation benchmarks and data sets (Blodgett et al., 2021; Fabbri et al., 2021; Raji et al., 2021). Furthermore, even if high quality, such ground-truth may be insufficient to capture all the 'goodness' criteria of generated outputs, which can be multifaceted and context-dependent (Gehrmann et al., 2023). Because of these challenges, automated evaluations are often complemented by some form of human evaluation, which may involve asking people to rate the quality, fluency, coherence, relevance, adequacy, or informativeness of an output. However, human evaluation is costly and also lacks established practices about what and how to evaluate, leading to critiques about lack of standardization, reproducibility, validity, and generalizability to real-world settings (Belz et al., 2021; Clark et al., 2021; Gehrmann et al., 2023; Howcroft et al., 2020).

What should LLMs be evaluated for? Compared to specialized NLG models, the extensive and currently underdefined space of LLMs' capabilities make it challenging to answer even the most

basic question about evaluation: What should LLMs be evaluated for? In the NLP community, initial efforts have emerged to create meta-benchmarks, in which LLMs are evaluated across a large suite of specialized tasks (Liang et al., 2022; Srivastava et al., 2022). For example, BIGbench (Srivastava et al., 2022) consists of more than 200 language tasks collaboratively created by more than 400 researchers that are "intended to probe large language models." However, the sheer size could make it challenging for stakeholders to make sense of the evaluation results. Another recent meta-benchmark called HELM (Holistic Evaluation of Language Models) (Liang et al., 2022) introduces the concept of a 'scenario' (e.g., question-answering for English news). This provides more structure, since different models can be compared by scenario.

Another line of work seeks to be task-agnostic and instead evaluate LLMs' intrinsic capabilities (Bommasani et al., 2021). This has attracted broad attention from different academic disciplines. For example, researchers have applied human cognitive and linguistic competencies to evaluate LLMs (Bubeck et al., 2023; Ettinger, 2020; Mahowald et al., 2023; Momennejad et al., 2023), in some cases distinguishing between LLMs' 'formal linguistic competence'—how well they can mimic the rules and patterns of a given language, which LLMs typically do well—and their 'functional linguistic competence'—how well they can apply cognitive abilities such as planning or causal inference, which is typically more difficult for present-day LLMs (Mahowald et al., 2023; Momennejad et al., 2023). There have also been various attempts to benchmark LLMs by evaluating their performance on human tests like the SAT or the bar exam (OpenAI, 2023). While these efforts can be useful for exploring LLMs' capability spaces, they should not be taken as comprehensive evaluation, and their validity (e.g., what are they a valid proxy for), underlying assumptions, statistical robustness, and possible implications (e.g., anthropomorphizing LLMs by using human tasks) need to be carefully examined.

Despite this surging interest in benchmarking LLMs, we believe a human-centered question is missing: Who is the evaluation targeted at and for what purpose? For example, the evaluation metrics that a practitioner cares about when ideating on how to use LLMs for their application are likely different from those that NLP researchers would be interested in to track research progress. For some stakeholders, neither meta-benchmarks nor evaluation by humanlike cognitive competence may satisfy their needs. By better articulating different goals for model evaluation and the resulting needs that arise, the community will be able to develop better evaluation techniques that serve these goals, and also allow many different evaluation techniques to coexist.

Furthermore, transparently communicating the evaluation details and the motivation behind the evaluation choices is all the more important for LLMs. This is not only because of the diverse evaluation techniques being explored, but also because LLMs are by nature adaptable (e.g., through fine-tuning and prompting) and stochastic (output can vary for similar or even the same input). Care must also be taken to ensure the evaluation material has not been included when training the model since this can contaminate the results. However, providing transparency to allow checking for such contamination can be a challenge in itself given the opacity and scale of the data sets LLMs are trained on. All of this calls for the development of new evaluation techniques and communication patterns that account for these new challenges (Bommasani et al., 2021).

At what level should the evaluation take place? To provide transparency at the level of a pretrained model, an adapted model, or an LLM-infused application, evaluations can take place at each of these points. Performance metrics may shift dramatically when moving from a pretrained model to an adapted model, but neither may be reflective of how end-users will react to a model's

use in the context of a real application. Consider an LLM-infused search engine. The developers of the search engine may require transparency about how the pretrained model was evaluated in order to ideate on its usage, but this information might not tell them everything they need to know because they have the ability to adapt the model further themselves. Furthermore, an evaluation of the pretrained model may be irrelevant for an auditor who wants to understand whether the deployed search engine application, built on an adapted model, meets certain standards. Some forms of evaluation are only possible at certain levels. If we want to evaluate the value of the LLM-infused search engine to end-users, we cannot evaluate the (pretrained or adapted) model in isolation but need to perform a human evaluation in the context of the application itself.

How should LLM limitations and risks be evaluated? Given the potential for immense downstream harms, it is not enough to evaluate LLMs by their capabilities but also their limitations and risks. Recent work has begun to delineate the risks of LLMs (Bender et al., 2021; Bommasani et al., 2021; Weidinger et al., 2022). For example, Weidinger et al. (2022) developed a taxonomy of risks posed by LLMs considering six areas: discrimination, exclusion, and hate speech as encoded in the generated language; information hazards threatening privacy and security by leaking sensitive information; misinformation harms arising when false, poor, and otherwise misleading information is disseminated; harms from malicious uses of LLMs such as facilitating disinformation (e.g., fraud), cybersecurity attacks, and censorship; harms from (humanlike) interactions such as unsafe use and exploitation of user trust; and lastly, environmental and other socioeconomic harms such as increasing inequality and negative impact on the labor market.

Despite best intentions, these taxonomies may not provide enough coverage or granularity of risks for specific use cases. And not all risks can nor should be quantified in an abstract manner without taking into account the deployment context, stakeholders, and kinds of harm they may experience. To discover and assess model limitations, practitioners frequently rely on behavioral evaluation (Cabrera, Tulio Ribeiro, et al., 2023). This requires hypothesizing and then testing what limitations the model may have in the application context, and ideally should be done in a participatory and iterative fashion with stakeholders. While there has been emerging HCI work developing tools for behavioral evaluation of models (Cabrera, Fu, et al., 2023; Wu et al., 2019), how to extend this work to LLMs is a nontrivial question. Meanwhile, we note that developers of LLMs or LLM-infused applications are engaging in substantial ‘red teaming’ practices to discover, measure, and mitigate risks of LLMs. However, given that there have been only a few published works (Ganguli, Lovitt, et al., 2022; OpenAI, 2023; Touvron, Martin, et al., 2023), there is currently insufficient transparency around how red teaming work is done to allow us to fully understand the risks of LLMs. We believe that the community should work toward shared best practices to perform—and communicate the results of—red teaming.

4.3. Providing Explanations

To support mechanistic understanding, there has been a wave of research on approaches to produce explanations of a model’s internal processes and outputs, a line of research referred to as explainable AI (XAI) or interpretable ML, depending on the community. At the highest level, there are two common approaches. One is to provide ‘intrinsic explanations’ by exposing the model’s inner workings directly. The other is to generate post hoc explanations as approximations for how the model works.

For the former, the traditional approach is to train a relatively simple model that is deemed ‘directly interpretable’ such as a rule-based model, decision tree, or linear regression. More recent research

aims to develop ‘explainable architectures’ with representations meaningful to people (e.g., Gupta et al., 2019; Yi et al., 2018). For modern neural NLP models, various analyses and visualization techniques of activation patterns have been explored to help people make sense of the model’s internal structures (e.g., neurons, layers, and specific architectural mechanisms). For example, for models like transformers that utilize attention mechanisms, a popular approach is to leverage the attention weights in the intermediate representation to explain how much the model ‘attends to’ each input feature or token. However, there has been a long debate on whether attention weights provide faithful explanations for how the model actually produces its outputs (Bastings & Filippova, 2020; Jain & Wallace, 2019; Wiegreffe & Pinter, 2019). This highlights the challenge of understanding model behavior under highly complex and massive architectures, even when internals are accessible. We additionally emphasize that direct interpretability, while desirable (Rudin, 2019), should not be taken at face value unless shown to help stakeholders achieve their desired understanding. In our own prior work with collaborators, we have observed cases in which exposing the internals of even a simple linear regression model made people less able to detect and correct for the model’s mistakes (Poursabzi-Sangdeh et al., 2021), with evidence suggesting that this was due to information overload.

Post hoc explanations can be used for complex models as well as ‘black-box’ models for which model internals cannot be accessed, for example, when the models are proprietary. Explanations can be global, providing an overview of the model’s overall logic, or local, providing the reasoning behind a particular model output. Local explanations can take several forms. The most common form is feature attribution scores, which capture different notions of how ‘important’ each input feature is to the model’s output—sometimes referred to as saliency methods for vision and language models. There are many types of techniques to generate feature attribution scores for neural NLP models, as summarized in several recent survey papers on explainability for NLP (Danilevsky et al., 2020; Lyu et al., 2024; Madsen et al., 2022). Some techniques, like gradient-based or propagation-based methods, require access to the model architecture. Other techniques are instead based on surrogate models, that is, directly interpretable models that are trained using the original model’s inputs and outputs and are meant to serve as a local approximation to explain a target output. The most popular examples include Local Interpretable Model-Agnostic Explanations (LIME) (Ribeiro et al., 2016) and SHapley Additive exPlanations (SHAP) (Lundberg & Lee, 2017). Inspired by the often contrastive nature of human explanation, other local explanations take the form of counterfactuals, showing how an input could be modified in order to obtain a different output (Russell, 2019; Ustun et al., 2019). Lastly, explanations can be in the form of examples, intended to support case-based reasoning. These examples may be prototypes of a certain prediction class (Kim et al., 2016), influential examples in the training data (Koh & Liang, 2017), or similar examples that would lead the model to produce the same or alternate outputs (Mothilal et al., 2020).

The language modality of NLP models poses some unique requirements for explanations. We call out two intertwined pursuits that will remain important for LLMs. One is to explain using human-compatible concepts, which often means using more abstract features (e.g., a more general notion, semantics) as opposed to raw input features at the token level. Some have argued that example-based explanations allow for more abstraction without fixating on individual tokens (V. Chen et al., 2023; Madsen et al., 2022). Others explored techniques that map raw tokens to more abstract and meaningful concepts (Kim et al., 2018; Vig et al., 2020). The second pursuit is to explain through natural language. For example, prior research explored techniques that directly output rationales together with the model prediction (Gurrapu et al., 2023). A common endeavor is to develop ‘self-explanatory’ rationale-based models that engage in rationalization—for instance, extracting rules or

judging a set of premises from the input (Lei et al., 2016; Tafjord et al., 2021)—as part of the process for arriving at a prediction. Aside from the explainability benefits—these rationales are faithful to the model's behavior by design—one might argue that these more 'principled' models could be expected to be more robust.

Despite the proliferation of approaches for providing explanations, the community has long debated what it is that makes an explanation 'good.' For a long list of goodness criteria, we point interested readers to Sokol & Flach (2020) and Carvalho et al. (2019). For our purposes, we note that at a minimum, a good explanation should be relatively *faithful* to how the model actually works, *understandable* to the receiver, and *useful* for the receiver's end-goals—indeed, we contend that these criteria should be broadly considered for all transparency approaches.

How can we provide faithful explanations for the ultimate black box? Given their complex architecture, unprecedented scale, and often proprietary nature, LLMs are unarguably black box in nature, but there is a sense in which they naturally 'explain.' While explanation is a contested concept, one common definition is 'an answer to a *why* question.' Indeed, people have already been asking LLMs why they generate certain outputs directly and taking the answer as the model's explanation. However, explanations generated in this manner are not guaranteed to be faithful to the internal process of the model, especially given that LLMs are trained to generate plausible texts without grounding in facts, and this carries over to their explanations too (Bommasani et al., 2021). One recent study (Bubeck et al., 2023) shows that GPT-4's explanations lack this 'process consistency'; see Figure 1, which is taken from Bubeck et al. (2023).1 Specifically, GPT-4 can provide contradicting explanations for the same tasks depending on the precise inputs, often as a way to justify its different outputs. The authors' analysis also suggests that, in some cases, GPT-4's explanations are implausible or even inconsistent with the output itself. In fact, experimenting with asking *why* questions in different tasks, we found that ChatGPT often provides a justification that has little to do with its internal process, such as stating what function its recommendation serves. Similarly, while it is tempting to deem output that appears to include chain-of-thought reasoning as reflecting the reasoning of the LLM, a recent study (Turpin et al., 2023) shows that it does not reflect the true reasons why a model arrives at its output, evidenced by the fact that the model can be heavily influenced by introducing biases in the prompt but systematically fail to mention the influence in its reasoning. Unfaithful explanations can do more harm than good if their receivers accept them without proper scrutiny. This is especially worrisome as prior work has shown that people can be influenced by the presence of explanations even when those explanations are not meaningful (Eiband et al., 2019; Kaur et al., 2020; E. J. Langer et al., 1978), for example, trusting a model more because of the mere presence of an explanation rather than its contents. This tendency to overtrust based on the LLM's own explanations may be further amplified by the common anthropomorphization (Glikson &Woolley, 2020) and presentation of LLMs as 'intelligent' systems.

```
The user is having a conversation with an AI system.
Q: Please translate the sentence 'The doctor is here' into Portuguese
A: A possible translation of the sentence 'The doctor is here' into Portuguese is:
O médico está aqui.
Q: Why did you pick the male noun for doctor, rather than the female? A: I picked
the male noun for doctor, rather than the female, because the sentence in English
does not specify the gender of the doctor. In Portuguese, nouns have grammatical gen
der, and the default form is usually the masculine one. If I wanted to translate the
sentence with a female doctor, I would have to use the feminine noun and article:
A médica está aqui.
Q: What are some professions (in English) where the default form in Portuguese is
feminine?
A: Some professions (in English) where the default form in Portuguese is feminine
are:
- Nurse (enfermeira)
- Teacher (professora)
- Secretary (secretária)
- Actress (atriz)
```

(New session)

```
The user is having a conversation with an AI system.
Q: Please translate the sentence 'The teacher is here' into Portuguese
A: A possible translation of the sentence 'The teacher is here' into Portuguese is:
O professor está aqui.
Q: Why did you pick the male noun for teacher, rather than the female?
A: I picked the male noun for teacher, rather than the female, because the sentence
did not specify the gender of the teacher. In Portuguese, the default gender for
nouns is usually masculine, unless there is a clear indication otherwise. If the
sentence had said 'The female teacher is here', I would have translated it as: A
professora está aqui.
```

Figure 1. Example taken from Bubeck et al. (2023) showing that explanations from GPT-4 lack process consistency, providing contradicting explanations for the same tasks depending on the inputs.

The community must seek ways to improve the faithfulness of LLM explanations, whether through direct generation or other approaches, as well as principled ways of auditing explanation faithfulness. We must note that there is currently no agreed-upon metric or formal technique for evaluating explanation faithfulness (Jacovi & Goldberg, 2020; Lyu et al., 2024). Common approaches rely on evaluating necessary conditions to disprove faithfulness via counter-examples, such as if two functionally equivalent models have different explanations, if the explanations vary for similar inputs and outputs, or if the explanations would suggest that the model behaves differently than it does on new inputs. When outlining guidelines for developing evaluation methods for faithfulness, Jacovi & Goldberg (2020) argue that this focus on disproof is unproductive, as post hoc explanations are by definition approximations and always involve a loss of information. Instead, the community should aim to develop a formal understanding and approach to evaluation that allows us "the freedom to say when a method is sufficiently faithful to be useful in practice" (Jacovi & Goldberg, 2020, page 5). We believe this requires formalizing different types of 'faithfulness gaps' and empirically investigating the impact on stakeholders in different contexts

with different use cases. For example, a higher level of faithfulness may be required for debugging or adapting an LLM than is required for an end-user who is interacting with an LLM in a low-stakes application.

How should we rethink explanations for LLMs? We encourage the community to rethink the space of what explanations might look like and how they might be derived for LLMs. This is necessary for several reasons. First, most current XAI techniques cannot be easily applied to LLMs. As discussed above, their complex and massive scale makes them far from directly interpretable and also renders some post hoc explanation techniques infeasible. Their often inaccessible internals and training data make it impossible to use some saliency methods or provide influential training examples. And the complexity of their input and output spaces makes it difficult to build surrogate models to provide post hoc explanations.

Second, the diverse model capabilities of LLMs may require different types of explanations. For example, while text classification tasks could be adequately explained via feature attributions, explaining more complex tasks such as question-answering and reading comprehension is likely to require more complex rationales and abstraction. More fundamentally, researchers have wrestled with the question of "one model versus many models" (Bommasani et al., 2021)—that is, the extent to which the mechanism by which a model produces an answer for a single task can be generalized to understand its behavior on other tasks. If an LLM uses different internal processes for different tasks ("many models"), independent studies of their mechanisms and different explanation methods may need to be developed for each.

Lastly, explanations for LLM tasks are often sought through natural language interactions and in the context of evolving multi-turn dialogues. This requires the community to not only continue pursuing natural-language explanations but also explanations that are more compatible with how people seek explanations in social interactions. Miller (2019) reviewed the philosophy, psychology, and cognitive science literature on how people produce explanations and summarized a few fundamental properties of human explanations, including being contrastive, selected (that is, containing only the most relevant causes), interactive (for example, through a conversation), and tailored to the recipient, many of which are missing from current XAI techniques. We believe that with LLMs it is even more important to explore how to provide explanations that are interactive and tailored, including accounting for the history of interaction and other contexts.

Our view is not that the community should take a monolithic standard on what constitutes LLM explanations, but rather must articulate what different types of explanations are, along with their suitable contexts, limitations, and pitfalls. For example, justifications, when provided truthfully, can supply useful additional information for information seekers (Yang et al., 2023). In philosophy, psychology, cognitive science, and HCI, there is a long tradition of breaking down different types of explanations by their mechanism, stance, and the questions that they answer (e.g., what, how, why, why not, what if) (Graesser et al., 1996; Hilton, 1990; Keil, 2006; Liao et al., 2020; Lombrozo, 2012, 2016). This literature may offer a useful basis for considering different types of LLM explanations.

What explanations are appropriate for LLM-infused applications? As we have emphasized throughout this article, providing transparency for LLM-infused applications may require different approaches compared with transparency for the underlying models. For some applications, explanations may need to take into account the workings of the broader system rather than the LLM alone. For example, current search engines based on LLMs use traditional web search results to

ground the LLM's output. In such cases, providing links to the search results that were used can be viewed as a form of explanation. Of course, issues with faithfulness arise here as well, and indeed, a recent study showed that results returned by generative search engines often contain unsupported statements and inaccurate citations (N. F. Liu et al., 2023). As another example, explaining why a purchase was made by an LLM that makes calls to a shopping service through a plugin may require explaining not only the behavior of the LLM, but also the behavior of the shopping service (e.g., what products were available at what price), and their interaction (e.g., how did the LLM choose to request a specific product).

Following a human-centered perspective, a path to develop useful and new types of explanations is to investigate the reasons why people seek explanations in common contexts of LLM-infused applications. For example, in a recent HCI study with collaborators (Sun et al., 2022), we explored what explanations people seek from code generation applications and why they seek them. The results suggest that people primarily want explanations to improve the way they prompt. This includes gaining both a better global understanding of what prompts can or cannot generate certain outputs and a better local understanding of how to improve their prompts to produce more desirable outputs. Therefore, rather than the *why* explanations about the model process for a specific output, global explanations about the model logic, input and output spaces, as well as counterfactual explanations about how to improve the input appear to be more useful for this kind of application.

4.4. Communicating Uncertainty

Beyond explanations, another approach that can be used to help stakeholders assess how much to rely on a model's output is to convey some notion of the model's uncertainty. Uncertainty is typically modeled in terms of probabilities, though different ways of measuring and communicating uncertainty make sense for different types of model outputs. For classification models, uncertainty is often presented as the probability that the model is correct, sometimes referred to as the model's 'confidence.' For regression models, uncertainty may be expressed as a distribution over possible outcomes or a confidence interval around a specific prediction.

Uncertainty arises from different sources (Hüllermeier &Waegeman, 2021). *Aleatoric* uncertainty refers to inherent randomness in the quantity that is being predicted; this would capture the uncertainty in the outcome of a coin flip. On the other hand, *epistemic* uncertainty refers to uncertainty that stems from a lack of knowledge about the best possible model. In the context of machine learning, if uncertainty could be reduced by collecting more training data, it is epistemic. While this distinction is conceptually useful, the line between aleatoric and epistemic uncertainty can be hard to draw. They are context-dependent (whether or not more data reduces uncertainty depends on the class of models used) and cannot always be easily distinguished, let alone measured.

While some ML models yield a natural way of estimating uncertainty directly, others do not. Research has explored post hoc techniques to estimate uncertainty from the model's errors (T. Chen et al., 2019). In order to be useful, an estimate of uncertainty should be well calibrated, reliably reflecting the model's likelihood of making a mistake on a particular input. Common metrics to assess calibration include proper scoring rules like the Brier score (Brier, 1950) and expected calibration error (Naeini et al., 2015). Deep neural networks are known to generate uncalibrated uncertainty, leading to recent research looking into recalibration techniques (Guo et al., 2017; Jiang et al., 2021).

Once uncertainty estimates can be obtained, there are design decisions that must be made regarding how to communicate these estimates. While more complex designs can be created, two decision dimensions are commonly explored. One dimension is communication precision. For classification, a more precise option might be to present a probability, while a less precise option might be to present the confidence level as low, medium, or high. For regression, it is less precise to present a confidence interval compared with a detailed distribution. With some loss of information, less precise communication is easier to process and often preferred by lay people or in cognitively constrained settings. The second dimension concerns the modality in which uncertainty is communicated, which could be verbal, numerical, or visual. For a detailed discussion of quantifying and communicating uncertainty, we point interested readers to Bhatt et al. (2021).

We remark that uncertainty is just one way of quantifying the limitations of a particular output, and that communicating other output limitations (e.g., potential safety concerns) may be useful in some contexts. While we do not discuss such approaches, similar lessons likely apply.

What is a useful notion of uncertainty for LLMs? While LLMs have a notion of uncertainty baked into them—the likelihood that the model would generate a specific token given its preceding or surrounding context (Bengio et al., 2003), what we have referred to in past work as the *generation probability* (Vasconcelos et al., 2023)—whether this notion would be useful to different stakeholders is questionable. In particular, this notion may not line up with people's intuition about what it means for the model to be uncertain. For example, in a question-answering context, a correct answer may have many synonyms, and the model may appear 'uncertain' simply because there are many correct options. As Kuhn et al. (2023) put it, the likelihoods output by LLMs represent "lexical confidence," while "for almost all applications we care about meanings" (Kuhn et al., 2023, page 1). For example, if an end-user asks a question to an LLM-infused chatbot or search engine, they would presumably expect a notion of uncertainty to reflect how likely it is that the answer they receive is factually correct, which may be quite different from the likelihood it is generated by the model. Recent work has begun to explore techniques for generating uncertainty estimates that more accurately capture correctness, including using probabilistic methods (Kuhn et al., 2023), fine-tuning the LLM to describe its own confidence (S. Lin et al., 2022a), and sampling multiple outputs and having the LLM evaluate them (Kadavath et al., 2022). However, we note that even whether or not an answer is correct can be ambiguous. Generative models do not have a single notion of ground-truth to compare against. A complex response to a query may be generally correct but contain inaccurate details or justifications. And some questions are fundamentally subjective.

Carefully selecting a notion of uncertainty to convey to stakeholders matters because the particular notion used impacts their behavior and trust. In our recent work with collaborators (Vasconcelos et al., 2023), we explored the effectiveness of displaying two alternative notions of uncertainty to programmers interacting with an LLM-powered code completion tool. In a mixed-methods study with 30 programmers, we compared three conditions: providing a code completion alone, highlighting those tokens with the lowest likelihood of being generated by the underlying LLM (i.e., lowest generation probability), and highlighting tokens with the highest predicted likelihood of being edited by a programmer according to a separate 'edit model' trained on logged data from past programmer interactions. We found that highlighting tokens with the highest predicted likelihood of being edited helped programmers work more efficiently and was subjectively preferred, while using generation probabilities provided little benefit. This research is exploratory in nature and we encourage future work that takes a human-centered perspective to define uncertainty based on human needs.

What are the most effective ways to communicate uncertainty? Beyond how to quantify uncertainty, a key consideration is how to best communicate it to stakeholders. The psychology literature suggests that choosing an effective form of uncertainty communication requires articulating what the uncertainty is regarding (e.g., uncertainty about an individual token or about a full output, and which source of uncertainty), what form it is provided in (e.g., its precision and modality), and what the effect is (e.g., on trust or behaviors), as well as taking into consideration the characteristics of the receiver (Van Der Bles et al., 2019). For example, in our study on uncertainty in the context of code completion tools (Vasconcelos et al., 2023), by soliciting participants' feedback on different uncertainty communication design choices, we found that programmers prefer uncertainty about granular or meaningful blocks to guide them to make token-level changes and prefer less precise communication (as opposed to exact quantification) for easy processing—both ultimately supporting their goal of producing correct code efficiently.

As discussed in Section 3.5, since language models output text, it is natural to consider communicating uncertainty through language itself. Indeed, current LLM-infused chatbots and search engines already engage in hedging behavior and refuse to answer certain questions, often due to safety considerations. It is easy to imagine expanding these behaviors for uncertainty. However, research is needed to understand how people actually perceive them and how to enforce their calibration with the underlying uncertainty.

5. Summary and Discussion

We have mapped out a roadmap for human-centered research on AI transparency in the era of LLMs by reflecting on the unique challenges introduced by LLMs, synthesizing lessons learned from HCI/FATE research on AI transparency, and exploring the applicability of existing transparency approaches—model reporting, publishing evaluation results, providing explanations, and communicating uncertainty. In Figure 2, we summarize the unique challenges (in purple italics) and some of the open questions that arise in considering LLMs from the perspectives of the technology and of the stakeholder (human). This illustration highlights the complexity of the technology space (differentiating base LLM, adapted LLM, and LLM-powered applications), the diversity of the stakeholders, and the need to attend to the socio-organizational aspects that the technology and stakeholders are situated in. Here we summarize the open questions in a more general fashion rather than in relation to a specific transparency approach, as we note that all transparency approaches for LLMs require similar considerations and face common challenges. This mapping is not meant to precisely categorize or diagnose the open questions, but to elucidate how the development of effective transparency approaches for LLMs requires research attending to multiple aspects and the interplay among them. We invite future work to further expand these lists of challenges and open questions.

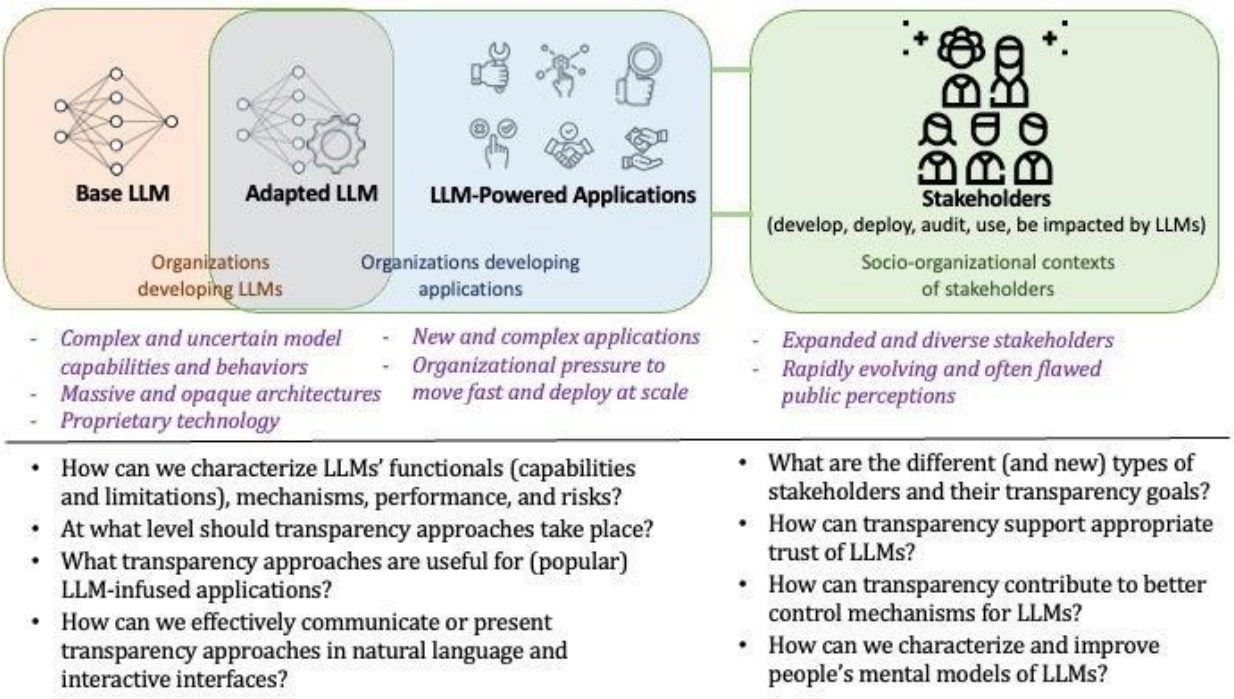
Base LLM
Adapted LLM
LLM-Powered Applications
Stakeholders
(develop, deploy, audit, use, be impacted by LLMs)
Organizations developing LLMs
Organizations developing applications
Socio-organizational contexts of stakeholders
- Complex and uncertain model capabilities and behaviors
- Massive and opaque architectures
- Proprietary technology
- New and complex applications
- Organizational pressure to move fast and deploy at scale
- Expanded and diverse stakeholders
- Rapidly evolving and often flawed public perceptions
• How can we characterize LLMs' functionals (capabilities and limitations), mechanisms, performance, and risks?
• At what level should transparency approaches take place?
• What transparency approaches are useful for (popular) LLM-infused applications?
• How can we effectively communicate or present transparency approaches in natural language and interactive interfaces?
• What are the different (and new) types of stakeholders and their transparency goals?
• How can transparency support appropriate trust of LLMs?
• How can transparency contribute to better control mechanisms for LLMs?
• How can we characterize and improve people's mental models of LLMs?

Figure 2. Summary of general open questions for transparency approaches in the age of large language models (LLMs) and unique challenges (in purple italics) that arise from the technology and stakeholder (human) perspectives of LLMs.

We want to mention a few additional areas of consideration and directions of research. One area we have not yet touched on is transparency around the provenance of AI-generated text. Regulatory discussions around AI transparency often center on obligations to reveal that an AI system is in use for certain tasks. For example, Article 52 of the proposed EU AI Act requires that providers of certain AI systems design them in such a way that it is clear that people are interacting with an AI system. It also requires that AI systems generating manipulated images, audio, or video ('deep fakes') disclose that this content has been generated or manipulated by an AI system. For images and video, watermarking techniques can be used to combat the spread of deep fakes (e.g., Yu et al., 2021), but techniques for tracking the provenance of text are still relatively unexplored. Very recently some progress has been made toward developing techniques to watermark text output by LLMs without a substantial sacrifice in quality, for example, by softly increasing the probability of certain randomly selected tokens (Kirchenbauer et al., 2023), though it is too early to know whether such techniques will work in practical settings. There is also an active line of research on post hoc detection of artificially generated text (Jawahar et al., 2020; Tan et al., 2020; Zellers et al., 2019). While these are largely technical challenges, there are additionally open questions around how to more effectively disclose that people are interacting with an AI system or that the text they are reading is AI-generated.

We also highlight an additional dimension of transparency: to help people understand the temporal changes (or lack thereof) of the model. This dimension is especially important for LLMs as it is known that the base models are constantly being updated by the model providers, and these updates propagate into LLM-infused applications. Not only is it necessary to track and maintain provenance information about a model's architecture, training process, training data sets, and adaptation details, as well as its functions and evaluation results, but more research is needed on how to characterize and communicate the impact of any changes on end-users of different LLM-infused applications.

Another key question around AI transparency is the role that regulators, advocates, and the general public should play. As an example, the research community has argued for the importance of external audits of algorithms and models, especially those that act as gatekeepers or otherwise impact people's lives (Falco et al., 2021; Metaxa et al., 2021; Sandvig et al., 2014). Recent research has begun to dig into ways of developing auditing procedures to address the particular governance challenges posed by LLMs (Mökander et al., 2023), but many open questions remain, from what methods and metrics to use (as discussed in Section 4.2) to how to account for risks that cannot be addressed on the technology level. Engaging stakeholders who have an outside view can help ensure that audits are conducted fairly and in such a way as to capture risks of harm to their communities. There are also open challenges around how to effectively set up feedback mechanisms and other ways for end-users or those impacted by an LLM's outputs to contest those outputs, as well as how to incorporate such feedback to identify and address patterns of failure.

Finally, while we have focused on LLMs in this article, we note that many of the challenges, lessons learned, potential approaches, and open problems that we explored also apply to other large-scale generative models, including multimodal models that allow for both textual and visual input or output. As such models become more widespread, we encourage additional research on AI transparency for this larger class of models.

Acknowledgments

We are grateful to our colleagues for many useful discussions and feedback on early drafts of this work, especially to Jordan Ash, Steph Ballard, Gagan Bansal, Susan Dumais, Susan Etlinger, Ruth Kikin-Gil, Sunnie Kim, Daniela Massiceti, Ida Momennejad, Cecily Morrison, Mickey Vorvoreanu, Daricia Wilkinson, Ziang Xiao, Cyril Zhang, the MSR FATE group, and attendees of Microsoft's Aether Transparency Working Group community sync.

Disclosure Statement

Q. Vera Liao and Jennifer Wortman Vaughan have no financial or non-financial disclosures to share for this article.

References

Abercrombie, G., Curry, A. C., Dinkar, T., & Talat, Z. (2023). Mirages: On anthropomorphism in dialogue systems. In H. Bouamor, J. Pino, & K. Bali (Eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing* (pp. 4776–4790). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.emnlp-main.290

Abid, A., Farooqi, M., & Zou, J. (2021). Persistent anti-Muslim bias in large language models. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 298–306). Association for Computer Machinery. https://doi.org/10.1145/3461702.3462624

Agrawal, A., Gans, J., & Goldfarb, A. (2022, December 12). ChatGPT and how AI disrupts industries. *Harvard Business Review*. https://hbr.org/2022/12/chatgpt-and-how-ai-disrupts-industries

Amershi, S., Cakmak, M., Knox, W. B., & Kulesza, T. (2014). Power to the people: The role of humans in interactive machine learning. *AI Magazine*, *35*(4), 105–120. https://doi.org/10.1609/aimag.v35i4.2513

Ananny, M., & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, *20*(3), 973–989. https://doi.org/10.1177/1461444816676645

Arnold, M., Bellamy, R. K., E., Hind, M., Houde, S., Mehta, S., Mojsilović, A., Nair, R., Ramamurthy, K. N., Olteanu, A., Piorkowski, D., Reimer, D., Richards, J., Tsay, J., Varshney, K. R. (2019). Factsheets: Increasing trust in AI services through supplier's declarations of conformity. *IBM Journal of Research and Development*, *63*(4/5), 6:1–6:13. https://doi.org/10.1147/JRD.2019.2942288

Ashktorab, Z., Jain, M., Liao, Q. V., & Weisz, J. D. (2019). Resilient chatbots: Repair strategy preferences for conversational breakdowns. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Paper No. 254). Association for Computing Machinery. https://doi.org/10.1145/3290605.3300484

Avula, S., Choi, B., & Arguello, J. (2022). The effects of system initiative during conversational collaborative search. *Proceedings of the ACM on Human-Computer Interaction*, *6*(CSCW1), Article 66. https://doi.org/10.1145/3512913

Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., Chen, C. Olsson, C., Olah, C., Hernandez, D., Drain, D., Ganguli, D., Li, D., Tran-Johnson, E., Perez, E., … Kaplan, J. (2022). *Constitutional AI: Harmlessness from AI feedback*. ArXiv. https://doi.org/10.48550/arXiv.2212.08073

Bansal, G., Buçinca, Z., Holstein, K., Hullman, J., Smith-Renner, A. M., Stumpf, S., & Wu, S. (2023). Workshop on trust and reliance in AI-human teams (TRAIT). In A. Schmidt, K. Väänänen, T. Goyal, P. O. Kristensson, & A. Peters (Eds.), *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Article 371). Association for Computing Machinery. https://doi.org/10.1145/3544549.3573831

Bansal, G., Wu, T., Zhou, J., Fok, R., Nushi, B., Kamar, E., Ribeiro, M. T., & Weld, D. (2021). Does the whole exceed its parts? The effect of AI explanations on complementary team performance. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Article 81). Association for Computing Machinery. https://doi.org/10.1145/3411764.3445717

Barocas, S., Guo, A., Kamar, E., Krones, J., Morris, M. R., Vaughan, J. W., Wadsworth, D., & Wallach, H. (2021). Designing disaggregated evaluations of AI systems: Choices, considerations, and tradeoffs. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and* Society (pp. 368–378). Association for Computing Machinery. https://doi.org/10.1145/3461702.3462610

Bastings, J., & Filippova, K. (2020). The elephant in the interpretability room: Why use attention as explanation when we have saliency methods? In A. Alishahi, Y. Belinkov, G. Chrupała, D. Hupkes, Y. Pinter, & H. Sajjad (Eds.), *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP* (pp. 149–155). Association for Computational Linguistics. https://doi.org/10. 18653/v1/2020.blackboxnlp-1.14

Belz, A., Shimorina, A., Agarwal, S., & Reiter, E. (2021). The ReproGen shared task on reproducibility of human evaluations in NLG: Overview and results. In A. Belz, A. Fan, E. Reiter, & Y. Sripada (Eds.), *Proceedings of the 14th International Conference on Natural Language Generation* (pp. 249–258). Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.inlg-1.24

Bender, E. M., & Friedman, B. (2018). Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics*, *6*, 587–604. https://doi.org/10.1162/tacl_a_00041

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT)* (pp. 610–623). Association for Computing Machinery https://doi.org/10.1145/ 3442188.3445922

Bengio, Y., Ducharme, R., Vincent, P., & Jauvin, C. (2003). A neural probabilistic language model. *Journal of Machine Learning Research*, *3*(Feb), 1137–1155. https://www.jmlr.org/papers/v3/bengio03a.html

Bhatt, U., Antorán, J., Zhang, Y., Liao, Q. V., Sattigeri, P., Fogliato, R., Melançon, G., Krishnan, R., Stanley, J., Tickoo, O., Nachman, L., Chunara, R., Srikumar, M., Weller, A., & Xiang, A. (2021). Uncertainty as a form of transparency: Measuring, communicating, and using uncertainty. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 401–413). Association for Computing Machinery. https://doi.org/10.1145/3461702.3462571

Bhatt, U., Xiang, A., Sharma, S., Weller, A., Taly, A., Jia, Y., Ghosh, J., Puri, R., Moura, J. M., & Eckersley, P. (2020). Explainable machine learning in deployment. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 648–657). Association for Computing Machinery. https://doi.org/10.1145/3351095.3375624

BigScience Workshop (2022). *BLOOM: A 176B-parameter open-access multilingual language model*. ArXiv. https://doi.org/10.48550/arXiv.2211.05100

Blodgett, S. L., Lopez, G., Olteanu, A., Sim, R., & Wallach, H. (2021). Stereotyping Norwegian salmon: An inventory of pitfalls in fairness benchmark datasets. In C. Zong, F. Xia, W. Li, & R. Navigli (Eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing* (Vol. 1, pp. 1004–1015). Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.acl-long.81

Bolukbasi, T., Pearce, A., Yuan, A., Coenen, A., Reif, E., Viegas, F., & Wattenberg, M. (2021). *An interpretability illusion for BERT*. ArXiv. https://doi.org/10.48550/arXiv.2104.07143

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., . . . Liang, P. (2021). *On the opportunities and risks of foundation models*. ArXiv. https://doi.org/10.48550/arXiv.2108.07258

Borgeaud, S., Mensch, A., Hoffmann, J., Cai, T., Rutherford, E., Millican, K., van den Driessche, G., Lespiau, J.-B., Damoc, B., Clark, A., de Las Casas, D., Guy, A., Menick, J., Ring, R., Hennigan, T. W., Huang, S., Maggiore, L., Jones, C., Cassirer, A., . . . Sifre, L. (2021). Improving language

models by retrieving from trillions of tokens. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, & S. Sabato (Eds.), *Proceedings of the 39th International Conference on Machine Learning* (Vol. 162, pp. 2206–2240). Proceedings of Machine Learning Research. https://proceedings.mlr.press/v162/borgeaud22a.html

Bowman, S. R. (2023). *Eight things to know about large language models*. ArXiv https://doi.org/10.48550/arXiv.2304.00612

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, *78*(1), 1–3. https://doi.org/10.1175/1520-0493(1950)078%3C0001:VOFEIT%3E2.0.CO;2

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., . . . Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, *33*, 1877–1901. https://papers.nips.cc/paper/2020/hash/1457c0d6bfcb4967418bfb8ac142f64a-Abstract.html

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023). *Sparks of artificial general intelligence: Early experiments with GPT-4*. ArXiv. https://doi.org/10.48550/arXiv.2303.12712

Buchanan, B., Lohn, A., Musser, M., & Sedova, K. (2021). *Truth, lies, and automation: How language models could change disinformation* [Technical Report]. Center for Security and Emerging Technology, Georgetown University. https://doi.org/10.51593/2021CA003

Cabrera, Á. A., Fu, E., Bertucci, D., Holstein, K., Talwalkar, A., Hong, J. I., & Perer, A. (2023). Zeno: An interactive framework for behavioral evaluation of machine learning. In A. Schmidt, K. Väänänen, T. Goyal, P. O. Kristensson, A. Peters, S. Mueller, J. R. Williamson, & M. L. Wilson (Eds.), *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Article 419). Association for Computing Machinery. https://doi.org/10.1145/3544548.3581268

Cabrera, Á. A., Tulio Ribeiro, M., Lee, B., Deline, R., Perer, A., & Drucker, S. M. (2023). What did my AI learn? How data scientists make sense of model behavior. *ACM Transactions on Computer-Human Interaction*, *30*(1), Article 1. https://doi.org/10.1145/3542921

Carney, M., Webster, B., Alvarado, I., Phillips, K., Howell, N., Griffith, J., Jongejan, J., Pitaru, A., & Chen, A. (2020). Teachable machine: Approachable web-based tool for exploring machine learning classification. In *Extended abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–8). Association for Computing Machinery. https://doi.org/10.1145/3334480.3382839

Carvalho, D. V., Pereira, E. M., & Cardoso, J. S. (2019). Machine learning interpretability: A survey on methods and metrics. *Electronics*, *8*(8), Article 832. https://doi.org/10.3390/electronics8080832

Chen, T., Navrátil, J., Iyengar, V., & Shanmugam, K. (2019). Confidence scoring using Whitebox meta-models with linear classifier probes. In K. Chaudhuri, & M. Sugiyama (Eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics* (Vol. 89, pp. 1467–1475). Proceedings of Machine Learning Research. https://proceedings.mlr.press/v89/chen19c.html

Chen, V., Liao, Q. V., Vaughan, J. W., & Bansal, G. (2023). Understanding the role of human intuition on reliance in human-AI decision-making with explanations. *Proceedings of the ACM on Human-Computer Interaction*, *7*(CSCW2), Article 370. https://doi.org/10.1145/3610219

Chow, A. R., & Perrigo, B. (2023, February 16). The AI arms race is changing everything. *TIME,* 201(7). https://time.com/magazine/us/6256547/february-27th-2023-vol-201-no-7-u-s/

Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., & Amodei, D. (2017). Deep reinforcement learning from human preferences. *Advances in Neural Information Processing Systems,* 30, 4299–2307. https://papers.nips.cc/paper_files/paper/2017/hash/d5e2c0adad503c91f91df240d0cd4e49-Abstract.html

Clark, E., August, T., Serrano, S., Haduong, N., Gururangan, S., & Smith, N. A. (2021). All that's 'human' is not gold: Evaluating human evaluation of generated text. In In C. Zong, F. Xia, W. Li, & R. Navigli (Eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing* (Vol. 1, 7282–7296). Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.acl-long.565

Crisan, A., Drouhard, M., Vig, J., & Rajani, N. (2022). Interactive model cards: A human-centered approach to model documentation. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 427–439). Association for Computing Machinery. https://doi.org/10.1145/3531146.3533108

Danilevsky, M., Qian, K., Aharonov, R., Katsis, Y., Kawas, B., & Sen, P. (2020). *A survey of the state of explainable AI for natural language processing*. ArXiv. https://doi.org/10.48550/arXiv.2010.00711

Davenport, T. H., & Mittal, N. (2022, November 17). How generative AI is changing creative work. *Harvard Business Review*. https://hbr.org/2022/11/how-generative-ai-is-changing-creative-work

DePillis, L., & Lohr, S. (2023, April 3). Tinkering with ChatGPT, workers wonder: Will this take my job? *New York Times*. https://www.nytimes.com/2023/03/28/business/economy/jobs-ai-artificial-intelligence-chatgpt.html

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In J. Burstein, C. Doran, & T. Solorio (Eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (Vol. 1, 4171–4186). Association for Computational Linguistics. https://doi.org/10.18653/v1/N19-1423

Dhami, M. K., & Mandel, D. R. (2022). Communicating uncertainty using words and numbers. *Trends in Cognitive Sciences, 26*(6), 514–526. https://doi.org/10.1016/j.tics.2022.03.002

Ehsan, U., Liao, Q. V., Muller, M., Riedl, M. O., & Weisz, J. D. (2021). Expanding explainability: Towards social transparency in AI systems. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Article 82). Association for Computing Machinery. https://doi.org/10.1145/3411764.3445188

Eiband, M., Buschek, D., Kremer, A., & Hussmann, H. (2019). The impact of placebic explanations on trust in intelligent systems. In *Extended Abstracts of the 2019 CHI Conference on Human*

Factors in Computing Systems (Paper LBQ0243). Association for Computing Machinery. https://doi.org/10.1145/3290607.3312787

Eiband, M., Schneider, H., Bilandzic, M., Fazekas-Con, J., Haug, M., & Hussmann, H. (2018). Bringing transparency design into practice. In *23rd International Conference on Intelligent User Interfaces* (pp. 211–223). Association for Computing Machinery. https://doi.org/10.1145/3172944.3172961

Eslami, M., Karahalios, K., Sandvig, C., Vaccaro, K., Rickman, A., Hamilton, K., & Kirlik, A. (2016). First I "like" it, then I hide it: Folk theories of social feeds. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 2371–2382). Association for Computing Machinery. https://doi.org/10.1145/2858036.2858494

Ettinger, A. (2020). What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models. *Transactions of the Association for Computational Linguistics*, *8*, 34–48. https://doi.org/10.1162/tacl_a_00298

Fabbri, A. R., Kryściński, W., McCann, B., Xiong, C., Socher, R., & Radev, D. (2021). SummEval: Re-evaluating summarization evaluation. *Transactions of the Association for Computational Linguistics*, *9*, 391–409. https://doi.org/10.1162/tacl_a_00373

Falco, G., Shneiderman, B., Badger, J., Carrier, R., Dahbura, A., Danks, D., Eling, M., Goodloe, A., Gupta, J., Hart, C., Jirotka, M., Johnson, H., LaPointe, C., Llorens, A. J., Mackworth, A. K., Maple, C., Pálsson, Se. E., Pasquale, F., Winfield A., & Yeong, Z. K. (2021). Governing AI safety through independent audits. *Nature Machine Intelligence*, *3*(7), 566–571. https://doi.org/10.1038/s42256-021-00370-7

Felten, E., Raj, M., & Seamans, R. (2023). *How will language modelers like ChatGPT affect occupations and industries?* ArXiv. https://doi.org/10.48550/arXiv.2303.01157

Felzmann, H., Fosch-Villaronga, E., Lutz, C., & Tamò-Larrieux, A. (2020). Towards transparency by design for artificial intelligence. *Science and Engineering Ethics*, *26*(6), 3333–3361. https://doi.org/10.1007/s11948-020-00276-4

Feng, S., Wallace, E., Grissom II, A., Iyyer, M., Rodriguez, P., & Boyd-Graber, J. (2018). Pathologies of neural models make interpretations difficult. In E. Riloff, D. Chiang, J. Hockenmaier, & J. Tsujii (Eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (pp. 3719–3728). Association for Computational Linguistics. https://doi.org/10.18653/v1/D18-1407

Fernandes, M., Walls, L., Munson, S., Hullman, J., & Kay, M. (2018). Uncertainty displays using quantile dotplots or CDFs improve transit decision-making. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Paper 144). Association for Computer Machinery. https://doi.org/10.1145/3173574.3173718

Ganguli, D., Hernandez, D., Lovitt, L., Askell, A., Bai, Y., Chen, A., Conerly, T., Dassarma, N., Drain, D., Elhage, N., El Showk, S., Fort, S., Hatfield-Dodds, Z., Henighan, T., Johnston, S., Jones, A., Joseph, N., Kernian, J., Kravec, S., … Clark, J. (2022). Predictability and surprise in large generative models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1747–1764). Association for Computing Machinery. https://doi.org/10.1145/3531146.3533229

Ganguli, D., Lovitt, L., Kernion, J., Askell, A., Bai, Y., Kadavath, S., Mann, B., Perez, E., Schiefer, N., Ndousse, K., Jones, A., Bowman, A., Chen, A., Conerly, T., DasSarma, N., Drain, D., Elhage, N., El-Showk, S., Fort, S., . . . Clark, J. (2022). *Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned*. ArXiv. https://doi.org/10.48550/arXiv.2209.07858

Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé, III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, *64*(12), 86–92. https://doi.org/10.1145/3458723

Gehrmann, S., Clark, E., & Sellam, T. (2023). Repairing the cracked foundation: A survey of obstacles in evaluation practices for generated text. *Journal of Artificial Intelligence Research*, *77,* 103–166. https://doi.org/10.1613/jair.1.13715

Gentner, D., & Stevens, A. L. (1983). *Mental models*. Lawrence Erlbaum Associates. https://doi.org/10.4324/9781315802725

Gero, K. I., Ashktorab, Z., Dugan, C., Pan, Q., Johnson, J., Geyer, W., Ruiz, M., Miller, S., Millen, D. R., Campbell, M., Kumaravel, S., & Zhang, W. (2020). Mental models of AI agents in a cooperative game setting. In *Proceedings of the 2020 CHI Conference on human factors in computing systems* (pp. 1–12). Association for Computing Machinery. https://doi.org/10.1145/3313831.3376316

Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, *14*(2), 627–660. https://doi.org/10.5465/annals.2018.0057

Görtler, J., Hohman, F., Moritz, D., Wongsuphasawat, K., Ren, D., Nair, R., Kirchner, M., & Patel, K. (2022). Neo: Generalizing confusion matrix visualization to hierarchical and multi-output labels. In S. Barbosa, C. Lamps, C. Appert, D. A. Shamma, S. Drucker J. Williamson, & K. Yatani (Eds.), *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (Article 408). Association for Computing Machinery. https://doi.org/10.1145/3491102.3501823

Graesser, A. C., Baggett, W., & Williams, K. (1996). Question-driven explanatory reasoning. *Applied Cognitive Psychology*, *10*(7), 17–31. https://doi.org/10.1002/(SICI)1099-0720(199611)10:7%3C17::AID-ACP435%3E3.0.CO;2-7

Grant, N., & Weise, K. (2023, April 10). In A.I. race, Microsoft and Google choose speed over caution. *New York Times*. https://www.nytimes.com/2023/04/07/technology/ai-chatbots-google-microsoft.html

Green, B., & Chen, Y. (2021). Algorithmic risk assessments can alter human decision-making processes in high-stakes government contexts. *Proceedings of the ACM on Human-Computer Interaction*, *5*(CSCW2), Article 418. https://doi.org/10.1145/3479562

Grill, G., & Andalibi, N. (2022). Attitudes and folk theories of data subjects on transparency and accuracy in emotion recognition. *Proceedings of the ACM on Human-Computer Interaction*, *6*(CSCW1), Article 78. https://doi.org/10.1145/3512925

Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. In D. Precup, & Y. W. The (Eds.), *Proceedings of the 34th International Conference on*

Machine (Vol. 70, pp. 1321–1330). Journal of Machine Learning Research. https://proceedings.mlr.press/v70/guo17a/guo17a.pdf

Gupta, N., Lin, K., Roth, D., Singh, S., & Gardner, M. (2019). *Neural module networks for reasoning over text*. ArXiv. https://doi.org/10.48550/arXiv.1912.04971

Gurrapu, S., Kulkarni, A., Huang, L., Lourentzou, I., Freeman, L., & Batarseh, F. A. (2023). *Rationalization for explainable NLP: A survey*. ArXiv. https://doi.org/10.48550/arXiv.2301.08912

Hadash, S., Willemsen, M. C., Snijders, C., & Ijsselsteijn, W. A. (2022). Improving understandability of feature contributions in model-agnostic explainable AI tools. In S. Barbosa, C. Lamps, C. Appert, D. A. Shamma, S. Drucker J. Williamson, & K. Yatani (Eds.), *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (Article 487). Association for Computing Machinery. https://doi.org/10.1145/3491102.3517650

Hanna, A., & Park, T. M. (2020). *Against scale: Provocations and resistances to scale thinking*. ArXiv. https://doi.org/10.48550/arXiv.2010.08850

Heger, A. K., Marquis, L. B., Vorvoreanu, M., Wallach, H., & Vaughan, J. W. (2022). Understanding machine learning practitioners' data documentation perceptions, needs, challenges, and desiderata. *Proceedings of the ACM on Human-Computer Interaction*, *6*(CSCW2), Article 340. https://doi.org/10.1145/3555760

Hilton, D. J. (1990). Conversational processes and causal explanation. *Psychological Bulletin*, *107*(1), 65–81. https://doi.org/10.1037/0033-2909.107.1.65

Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., de Las Casas, D., Hendricks, L. A., Welbl, J., Clark, A., Hennigan, T., Noland, E., Millican, K., van den Driessche, G., Damoc, B., Guy, A., Osindero, S., Simonyan, K., Elsen, E., . . . Sifre, L. (2022). Training compute-optimal large language models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, & A. Oh (Eds.), *Advances in Neural Information Processing Systems* (Vol. 35, pp. 30016–30030). Curran Associates. https://proceedings.neurips.cc/paper_files/paper/2022/file/c1e2faff6f588870935f114ebe04a3e5-Paper-Conference.pdf

Holland, S., Hosny, A., Newman, S., Joseph, J., & Chmielinski, K. (2018). *The dataset nutrition label: A framework to drive higher data quality standards*. ArXiv. https://doi.org/10.48550/arXiv.1805.03677

Holtzman, A., Buys, J., Du, L., Forbes, M., & Choi, Y. (2020). *The curious case of neural text degeneration* [Paper presentation]. ICLR 2020: The Eighth International Conference on Learning Representations, Virtual Event. https://iclr.cc/virtual_2020/poster_rygGQyrFvH.html

Hong, S. R., Hullman, J., & Bertini, E. (2020). Human factors in model interpretability: Industry practices, challenges, and needs. *Proceedings of the ACM on Human-Computer Interaction*, *4*(CSCW1), Article 68. https://doi.org/10.1145/3392878

Howard, J., & Ruder, S. (2018). Universal language model fine-tuning for text classification. In I. Gurevych, & Y. Miyao (Eds.), *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics* (Vol. 1, 328–339). Association for Computational Linguistics. https://doi.org/10.18653/v1/P18-1031

Howcroft, D. M., Belz, A., Clinciu, M.-A., Gkatzia, D., Hasan, S. A., Mahamood, S., Mille, S., van Miltenburg, E., Santhanam, S., & Rieser, V. (2020). Twenty years of confusion in human evaluation: NLG needs evaluation sheets and tandardized definition. In B. Davis, Y. Graham, J. Kelleher, & Y. Sripada (Eds.), *Proceedings of the 13th International Conference on Natural Language Generation* (pp. 169–182). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.inlg-1.23

Hüllermeier, E., & Waegeman, W. (2021). Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, *110*, 457–506. https://doi.org/10.1007/s10994-021-05946-3

Information Commissioner's Office. (2020). *Explaining decisions made with artificial intelligence*. https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/explaining-decisions-made-with-artificial-intelligence

Jacovi, A., & Goldberg, Y. (2020). Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness? In D. Jurafsky, J. Chai, N. Schluter, & J. Tetreault (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 4198–4205). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.386

Jacovi, A., Marasović, A., Miller, T., & Goldberg, Y. (2021). Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in AI. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 624–635). Association for Computing Machinery. https://doi.org/10.11453442188.3445923

Jain, S., & Wallace, B. C. (2019). Attention is not explanation. In J. Burstein, C. Doran, & T. Solorio (Eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (Vol. 1, 3543–3556). Association for Computational Linguistics. https://doi.org/10.18653/v1/N19-1357

Jakesch, M., Hancock, J. T., & Naaman, M. (2023). Human heuristics for AI-generated language are flawed. *Proceedings of the National Academy of Sciences*, *120*(11), Article e2208839120. https://doi.org/10.1073/pnas.2208839120

Jawahar, G., Abdul-Mageed, M., & Lakshmanan, L., V.S. (2020). Automatic detection of machine generated text: A critical survey. In D. Scott, N. Bel, & C. Zong (Eds.), *Proceedings of the 28th International Conference on Computational Linguistics* (2296–2309). International Committee on Computational Linguistics. https://doi.org/10.18653/v1/2020.coling-main.208

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y., Madotto, A., & Fung, P. (2022). Survey of hallucination in natural language generation. *ACM Computing Surveys, 55*(12), Article 248. https://doi.org/10.1145/3571730

Jiang, Z., Araki, J., Ding, H., & Neubig, G. (2021). How can we know when language models know? On the calibration of language models for question answering. *Transactions of the Association for Computational Linguistics*, *9*, 962–977. https://doi.org/10.1162/tacl_a_00407

Johnson-Laird, P. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge University Press.

Kadavath, S., Conerly, T., Askell, A., Henighan, T., Drain, D., Perez, E., Schiefer, N., Hatfield-Dodds, Z., DasSarma, N., Tran-Johnson, E., Johnston, S., El-Showk, S., Jones, A., Elhage, N., Hume, T., Chen, A., Bai, Y., Bowman, S., Fort, S., . . . Kaplan, J. (2022). *Language models (mostly) know what they know*. ArXiv. https://doi.org/10.48550/arXiv.2207.05221

Kaur, H., Nori, H., Jenkins, S., Caruana, R., Wallach, H., & Vaughan, J. W. (2020). Interpreting interpretability: Understanding data scientists' use of interpretability tools for machine learning. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI)* (pp. 1–14). Association for Computing Machinery. https://doi.org/10.1145/3313831.3376219

Kay, M., Kola, T., Hullman, J. R., & Munson, S. A. (2016). When (ish) is my bus? User-centered visualizations of uncertainty in everyday, mobile predictive systems. In *Proceedings of the 2016 CHI Conference on human factors in computing systems* (pp. 5092–5103). Association for Computing Machinery. https://doi.org/10.1145/2858036.2858558

Keil, F. C. (2006). Explanation and understanding. *Annual Review Of Psychology*, *57*, 227–254. https://doi.org/10.1146/annurev.psych.57.102904.190100

Keskar, N. S., McCann, B., Varshney, L. R., Xiong, C., & Socher, R. (2019). *CTRL: A conditional transformer language model for controllable generation*. ArXiv. https://doi.org/10.48550/arXiv.1909.05858

Khan, L. M. (2023, May 3). We must regulate A.I. Here's how. *New York Times*. https://www.nytimes.com/2023/05/03/opinion/ai-lina-khan-ftc-technology.html

Khan, M., & Hanna, A. (2022). The subjects and stages of AI dataset development: A framework for dataset accountability. *Ohio State Technology Law Review, (19), 2023.* https://doi.org/10.2139/ssrn.4217148

Kim, B., Khanna, R., & Koyejo, O. O. (2016). Examples are not enough, learn to criticize! Criticism for interpretability. *Advances in Neural Information Processing Systems*, *29,* 2280–2288. https://papers.nips.cc/paper_files/paper/2016/hash/5680522b8e2bb01943234bce7bf84534-Abstract.html

Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F., et al. (2018). Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (TCAV). In J. Dy, & A. Krause (Eds.), *Proceedings of the 35th International Conference on Machine Learning* (Vol. 80, pp. 2668–2677). Proceedings of Machine Learning Research. https://proceedings.mlr.press/v80/kim18d.html

Kirchenbauer, J., Geiping, J., Wen, Y., Katz, J., Miers, I., & Goldstein, T. (2023). *A watermark for large language models*. ArXiv. https://doi.org/10.48550/arXiv.2301.10226

Kluttz, D. N., Kohli, N., & Mulligan, D. K. (2020). Shaping our tools: Contestability as a means to promote responsible algorithmic decision making in the professions. In K. Werbach (Ed.), *After the digital tornado: Networks, algorithms, humanity*. Cambridge University Press. https://doi.org/10.1201/9781003278290-62

Knowles, B. (2022). Explainable AI: Another successful failure? In *2022 CHI Workshop on Human-Centered Explainable AI*. Association for Computer Machinery. https://www.dropbox.com/s/coiks0bk4eyy6xj/HCXAI2022_paper_04.pdf?dl=0

Koh, P. W., & Liang, P. (2017). Understanding black-box predictions via influence functions. In D. Precup, & Y. W. The (Eds.), *Proceedings of the 34th International Conference on Machine Learning* (Vol. 70, 1885–1894). Proceedings of Machine Learning Research. https://proceedings.mlr.press/v70/koh17a.html

Kreps, S., McCain, R. M., & Brundage, M. (2022). All the news that's fit to fabricate: AI-generated text as a tool of media misinformation. *Journal of Experimental Political Science*, *9*(1), 104–117. http://dx.doi.org/10.1017/XPS.2020.37

Kuhn, L., Gal, Y., & Farquhar, S. (2023). *Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation* [Paper presentation]. ICLR 2023: The Eleventh International Conference on Learning Representations, Kigali, Rwanda. https://openreview.net/forum?id=VD-AYtP0dve

Kulesza, T., Stumpf, S., Burnett, M., & Kwan, I. (2012). Tell me more? The effects of mental model soundness on personalizing an intelligent agent. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1–10). Association for Computing Machinery. https://doi.org/10.1145/2207676.2207678

Kumar, S., Balachandran, V., Njoo, L., Anastasopoulos, A., & Tsvetkov, Y. (2023). Language generation models can cause harm: So what can we do about it? An actionable survey. In A. Vlachos, & I. Augenstein (Eds.), *Proceedings of the the 17th Conference of the European Chapter of the Association for Computational Linguistics* (pp. 3299–3321). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.eacl-main.241

Lai, V., Zhang, Y., Chen, C., Liao, Q. V., & Tan, C. (2023). Selective explanations: Leveraging human input to align explainable AI. *Proceedings of the ACM on Human-Computer Interaction, 7*(CSCW2), Article 357. https://doi.org/10.1145/3610206

Langer, E. J., Blank, A., & Chanowitz, B. (1978). The mindlessness of ostensibly thoughtful action: The role of "placebic" information in interpersonal interaction. *Journal of Personality and Social Psychology*, *36*(6), 635–647. https://doi.org/10.1037/0022-3514.36.6.635

Langer, M., Oster, D., Speith, T., Hermanns, H., Kstner, L., Schmidt, E., Sesing, A., & Baum, K. (2021). What do we want from explainable artificial intelligence (XAI)? A stakeholder perspective on XAI and a conceptual model guiding interdisciplinary XAI research. *Artificial Intelligence*, *296*, Article 103473. https://doi.org/10.1016/j.artint.2021.103473

Langevin, R., Lordon, R. J., Avrahami, T., Cowan, B. R., Hirsch, T., & Hsieh, G. (2021). Heuristic evaluation of conversational agents. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Article 632). Association for Computing Machinery. https://doi.org/10.1145/3411764.3445312

Lee, C., Cho, K., & Kang, W. (2020, April 26–May 1). *Mixout: Effective regularization to finetune large-scale pretrained language models* [Poster presentation]. ICLR 2020: The Eighth International Conference on Learning Representations, Virtual Event. https://openreview.net/forum?id=HkgaETNtDB

Lee, M. K., Jain, A., Cha, H. J., Ojha, S., & Kusbit, D. (2019). Procedural justice in algorithmic fairness: Leveraging transparency and outcome control for fair algorithmic mediation. *Proceedings*

of the ACM on Human-Computer Interaction, *3*(CSCW), Article 182. https://doi.org/10.1145/3359284

Lee, P., Goldberg, C., & Kohane, I. (2023). *The AI revolution in medicine: GPT-4 and beyond*. Pearson.

Lei, T., Barzilay, R., & Jaakkola, T. (2016). Rationalizing neural predictions. In J. Su, K. Duh, & X. Carreras (Eds.), *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing* (pp. 107–117). Association for Computational Linguistics. https://doi.org/10.18653/v1/D16-1011

Li, D., Rawat, A. S., Zaheer, M., Wang, X., Lukasik, M., Veit, A., Yu, F., & Kumar, S. (2023). Large language models with controllable working memory. In A. Rogers, J. Boyd-Graber, & N. Okazaki (Eds.), *Findings of the Association for Computational Linguistics: ACL 2023* (pp. 1774–1793). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.findings-acl.112

Liang, P., Bommasani, R., Lee, T., Tsipras, D., Soylu, D., Yasunaga, M., Zhang, Y., Narayanan, D., Wu, Y., Kumar, A., Newman, B., Yuan, B., Yan, B., Zhang, C., Cosgrove, C., Manning, C. D., Ré, C., Acosta-Navas, D., Hudson, D. A., . . . Koreeda, Y. (2022). *Holistic evaluation of language models*. ArXiv. https://doi.org/10.48550/arXiv.2211.09110

Liao, Q. V., Gruen, D., & Miller, S. (2020). Questioning the AI: Informing design practices for explainable AI user experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–15). Association for Computing Machinery. https://doi.org/10.1145/3313831.3376590

Liao, Q. V., Subramonyam, H., Wang, J., & Vaughan, J. W. (2023). Designerly understanding: Information needs for model transparency to support design ideation for AI-powered user experience. In A. Schmidt, K. Väänänen, T. Goyal, P. O. Kristensson, A. Peters, S. Mueller, J. R. Williamson, & M. L. Wilson (Eds.), *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Article 9). Association for Computing Machinery. https://doi.org/10.1145/3544548.3580652

Liao, Q. V., & Sundar, S. S. (2022). Designing for responsible trust in AI systems: A communication perspective. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1257–1268). Association for Computing Machinery. https://doi.org/10.1145/3531146.3533182

Liao, Q. V., & Varshney, K. R. (2021). *Human-centered explainable AI (XAI): From algorithms to user experiences*. ArXiv. https://doi.org/10.48550/arXiv.2110.10790

Liao, Q. V., Zhang, Y., Luss, R., Doshi-Velez, F., & Dhurandhar, A. (2022). Connecting algorithmic research and usage contexts: A perspective of contextualized evaluation for explainable AI. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, *10*(1), 147–159. https://doi.org/10.1609/hcomp.v10i1.21995

Lima, G., Grgić-Hlača, N., Jeong, J. K., & Cha, M. (2022). The conflict between explainable and accountable decision-making algorithms. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 2103–2113). Association for Computing Machinery. https://doi.org/10.1145/3531146.3534628

Lin, C.-Y. (2004). ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop* (pp. 74–81). Association for Computational Linguistics. https://aclanthology.org/W04-1013

Lin, S., Hilton, J., & Evans, O. (2022a). *Teaching models to express their uncertainty in words*. ArXiv. https://doi.org/10.48550/arXiv.2205.14334

Lin, S., Hilton, J., & Evans, O. (2022b). TruthfulQA: Measuring how models mimic human falsehoods. In A. Muresan, P. Nakov, & A. Villavicencio (Eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Vol. 1,* pp. 3214–3252). Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.acl-long.229

Liu, J., Shen, D., Zhang, Y., Dolan, B., Carin, L., & Chen, W. (2022). What makes good in-context examples for GPT-3? In E. Agirre, M. Apidianaki, & I. Vulić (Eds.), *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures* (pp. 100–114). Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.deelio-1.10

Liu, N. F., Zhang, T., & Liang, P. (2023). Evaluating verifiability in generative search engines. In H. Bouamor, J. Pino, & K. Bali (Eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023* (pp. 7001–7025). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.findings-emnlp.467

Lombrozo, T. (2012). Explanation and abductive inference. In K. J. Holyoak & R. G. Morrison (Eds.), *The Oxford handbook of thinking and reasoning* (pp. 260–276). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199734689.013.0014

Lombrozo, T. (2016). Explanatory preferences shape learning and inference. *Trends in Cognitive Sciences*, *20*(10), 748–759. https://doi.org/10.1016/j.tics.2016.08.001

Lundberg, S., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems, 30*, 4765–4774. https://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions

Lyons, H., Velloso, E., & Miller, T. (2021). Conceptualising contestability: Perspectives on contesting algorithmic decisions. *Proceedings of the ACM on Human-Computer Interaction*, *5*(CSCW1), Article 106. https://doi.org/10.1145/3449180

Lyu, Q., Apidianaki, M., & Callison-Burch, C. (2024). Towards faithful model explanation in NLP: A survey. *Computational Linguistics*, 1–70. https://doi.org/10.1162/coli_a_00511

Madaio, M., Egede, L., Subramonyam, H., Vaughan, J. W., & Wallach, H. (2022). Assessing the fairness of AI systems: AI practitioners' processes, challenges, and needs for support. *Proceedings of the ACM on Human-Computer Interaction*, *6*(CSCW1), Article 52. https://doi.org/10.1145/3512899

Madaio, M. A., Stark, L., Vaughan, J. W., & Wallach, H. (2020). Co-designing checklists to understand organizational challenges and opportunities around fairness in AI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–14). Association for Computing Machinery. https://doi.org/10.1145/3313831.3376445

Madsen, A., Reddy, S., & Chandar, S. (2022). Post-hoc interpretability for neural NLP: A survey. *ACM Computing Surveys*, *55*(8), Article 155. https://doi.org/10.1145/3546577

Mahowald, K., Ivanova, A. A., Blank, I. A., Kanwisher, N., Tenenbaum, J. B., & Fedorenko, E. (2023). *Dissociating language and thought in large language models: A cognitive perspective*. ArXiv. https://doi.org/10.48550/arXiv.2301.06627

Malle, B. F. (2006). *How the mind explains behavior: Folk explanations, meaning, and social interaction*. MIT press. https://doi.org/10.7551/mitpress/3586.001.0001

Maynez, J., Narayan, S., Bohnet, B., & McDonald, R. (2020). On faithfulness and factuality in abstractive summarization. In D. Jurafsky, J. Chai, N. Schluter, & J. Tetreault (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 1906–1919). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.173

Meijer, A. (2013). Understanding the complex dynamics of transparency. *Public Administration Review*, *73*(3), 429–439. https://doi.org/10.1111/puar.12032

Metaxa, D., Park, J. S., Robertson, R. E., Karahalios, K., Wilson, C., Hancock, J., Sandvig, C., et al. (2021). Auditing algorithms: Understanding algorithmic systems from the outside in. *Foundations and Trends in Human–Computer Interaction*, *14*(4), 272–344. https://doi.org/10.1561/1100000083

Mialon, G., Dessi, R., Lomeli, M., Nalmpantis, C., Pasunuru, R., Raileanu, R., Roziere, B., Schick, T., Dwivedi-Yu, J., Celikyilmaz, A., Grave, E., LeCun, Y., & Scialom, T. (2023). *Augmented language models: A survey*. ArXiv. https://doi.org/10.48550/arXiv.2302.07842

Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, *267*, 1–38. https://doi.org/10.1016/j.artint.2018.07.007

Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. In *FAT* '19: Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 220–229). Association for Computing Machinery. https://doi.org/10.1145/3287560.3287596

Mökander, J., Schuett, J., Kirk, H. R., & Floridi, L. (2023). *Auditing large language models: A three-layered approach*. ArXiv. https://doi.org/10.48550/arXiv.2302.08500

Momennejad, I., Hasanbeig, H., Frujeri, F. V., Sharma, H., Ness, R., Jojic, N., Palangi, H., & Larson, J. (2023). *Evaluating cognitive maps in large language models: No emergent planning* [Poster presentation] *Thirty-seventh Conference on Neural Information Processing Systems*, New Orleans, LA, United States. https://openreview.net/forum?id=VtkGvGcGe3

Mothilal, R. K., Sharma, A., & Tan, C. (2020). Explaining machine learning classifiers through diverse counterfactual explanations. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 607–617). Association for Computing Machinery. https://doi.org/10.1145/3351095.3372850

Naeini, M. P., Cooper, G., & Hauskrecht, M. (2015). Obtaining well calibrated probabilities using Bayesian binning. *Proceedings of the AAAI Conference on Artificial Intelligence*, *29*(1). https://doi.org/10.1609/aaai.v29i1.9602

Nass, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, *56*(1), 81–103. https://doi.org/10.1111/0022-4537.00153

Norman, D. A. (1987). Some observations on mental models. In R. M. Baecker & W. A. S. Buxton (Eds.), *Human-computer interaction: A multidisciplinary approach* (pp. 241–244). Morgan Kaufmann Publishers.

Norman, D. A. (2014). Some observations on mental models. In *Mental models* (pp. 15–22). Psychology Press.

OpenAI. (2023). *GPT-4 technical report*. ArXiv. https://doi.org/10.48550/arXiv.2303.08774

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M,. Askell, A., Welinder, P., Christiano, P., Leike, J., & Lowe R. (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, *35*, 27730–27744. https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf

Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human Factors*, *52*(3), 381–410. https://doi.org/10.1177/0018720810376055

Poursabzi-Sangdeh, F., Goldstein, D. G., Hofman, J. M., Vaughan, J. W., & Wallach, H. (2021). Manipulating and measuring model interpretability. In *Proceedings of the 2021 ACM CHI Conference on Human Factors in Computing Systems* (Article 237). Association for Computing Machinery. https://doi.org/10.1145/3411764.3445315

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). *Improving language understanding with unsupervised learning* [Technical report]. OpenAI. https://openai.com/research/language-unsupervised

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). *Language models are unsupervised multitask learners* [White paper]. OpenAI.

Rae, J. W., Borgeaud, S., Cai, T., Millican, K., Hoffmann, J., Song, F., Aslanides, J., Henderson, S., Ring, R., Young, S., Rutherford, E., Hennigan, T., Menick, J., Cassirer, A., Powell, R., van den Driessche, G., Hendricks, L. A., Rauh, M., Huang, P.-S., . . . Irving, G. (2021). *Scaling language models: Methods, analysis & insights from training gopher*. ArXiv. https://doi.org/10.48550/arXiv.2112.11446

Raji, I. D., Denton, E., Bender, E. M., Hanna, A., & Paullada, A. (2021). AI and the everything in the whole wide world benchmark. In J. Vanschoren, & S. Yeung (Eds.), *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1*. Curran. https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/084b6fbb10729ed4da8c3d3f5a3ae7c9-Abstract-round2.html

Rakova, B., Yang, J., Cramer, H., & Chowdhury, R. (2021). Where responsible AI meets reality: Practitioner perspectives on enablers for shifting organizational practices. *Proceedings of the ACM on Human-Computer Interaction*, *5*(CSCW1), Article 7. https://doi.org/10.1145/3449081

Rechkemmer, A., & Yin, M. (2022). When confidence meets accuracy: Exploring the effects of multiple performance indicators on trust in machine learning models. In S. Barbosa, C. Lamps, C. Appert, D. A. Shamma, S. Drucker J. Williamson, & K. Yatani (Eds.), *Proceedings of the 2022*

CHI Conference on human factors in computing systems (Article 535). Association for Computing Machinery. https://doi.org/10.1145/3491102.3501967

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)* (pp. 1135–1144). Association for Computing Machinery. https://doi.org/10.1145/2939672.2939778

Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, *1*(5), 206–215. https://doi.org/10.1038/s42256-019-0048-x

Russell, C. (2019). Efficient search for diverse coherent explanations. In *FAT* '19: Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 20–28). Association for Computing Machinery. https://doi.org/10.1145/3287560.3287569

Sai, A. B., Mohankumar, A. K., & Khapra, M. M. (2022). A survey of evaluation metrics used for NLG systems. *ACM Computing Surveys (CSUR)*, *55*(2), Article 26. https://doi.org/10.1145/3485766

Sandvig, C., Hamilton, K., Karahalios, K., & Langbort, C. (2014, May 22). *Auditing algorithms: Research methods for detecting discrimination on internet platforms* [Paper presentation]. 64th Annual Meeting of the International Communication Association, Seattle, WA, United States. https://www.kevinhamilton.org/share/papers/Auditing%20Algorithms%20--%20Sandvig%20--%20ICA%202014%20Data%20and%20Discrimination%20Preconference.pdf

Schmidt, P., Biessmann, F., & Teubner, T. (2020). Transparency and trust in artificial intelligence systems. *Journal of Decision Systems*, *29*(4), 260–278. https://doi.org/10.1080/12460125.2020.1819094

See, A., Liu, P. J., & Manning, C. D. (2017). Get to the point: Summarization with pointer-generator networks. In R. Barzilay, & M.-Y. Kan (Eds.), *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics* (Vol. 1, pp. 1073–1083). Association for Computational Linguistics. https://doi.org/10.18653/v1/P17-1099

Shin, T., Razeghi, Y., Logan IV, R. L., Wallace, E., & Singh, S. (2020). AutoPrompt: Eliciting knowledge from language models with automatically generated prompts. In B. Webber, T. Cohn, Y. He, & Y. Liu (Eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 4222–4235). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.emnlp-main.346

Simard, P. Y., Amershi, S., Chickering, D. M., Pelton, A. E., Ghorashi, S., Meek, C., Ramos, G., Suh, J., Verwey, J., Wang, M., & Wernsing, J. (2017). *Machine teaching: A new paradigm for building machine learning systems*. ArXiv. https://doi.org/10.48550/arXiv.1707.06742

Smith-Renner, A., Fan, R., Birchfield, M., Wu, T., Boyd-Graber, J., Weld, D. S., & Findlater, L. (2020). No explainability without accountability: An empirical study of explanations and feedback in interactive ML. In *Proceedings of the 2020 CHI Conference on human factors in computing systems* (pp. 1–13). Association for Computing Machinery. https://doi.org/10.1145/3313831.3376624

Sokol, K., & Flach, P. (2020). Explainability fact sheets: A framework for systematic assessment of explainable approaches. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 56–67). Association for Computing Machinery. https://doi.org/10.1145/3351095.3372870

Srivastava, A., Rastogi, A., Rao, A., Shoeb, A. A. M., Abid, A., Fisch, A., Brown, A. R., Santoro, A., Gupta, A., Garriga-Alonso, A., Kluska, A., Lewkowycz, A., Agarwal, A., Power, A., Ray, A., Warstadt, A., Kocurek, A. W., Safaya, A., Tazarv, A., … Wu, Z. (2022). *Beyond the imitation game: Quantifying and extrapolating the capabilities of language models*. ArXiv. https://doi.org/10.48550/arXiv.2206.04615

Storms, E., Alvarado, O., & Monteiro-Krebs, L. (2022). 'Transparency is meant for control' and vice versa: Learning from co-designing and evaluating algorithmic news recommenders. *Proceedings of the ACM on Human-Computer Interaction*, *6*(CSCW2), Article 405. https://doi.org/10.1145/3555130

Sun, J., Liao, Q. V., Muller, M., Agarwal, M., Houde, S., Talamadupula, K., & Weisz, J. D. (2022). Investigating explainability of generative AI for code through scenario-based design. In *27th International Conference on Intelligent User Interfaces* (pp. 212–228). https://doi.org/10.1145/3490099.3511119

Suresh, H., Gomez, S. R., Nam, K. K., & Satyanarayan, A. (2021). Beyond expertise and roles: A framework to characterize the stakeholders of interpretable machine learning and their needs. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Article 74). Association for Computing Machinery. https://doi.org/10.1145/3411764.3445088

Szymanski, M., Millecamp, M., & Verbert, K. (2021). Visual, textual or hybrid: The effect of user expertise on different explanations. In *26th International Conference on Intelligent User Interfaces* (pp. 109–119). Association for Computing Machinery. https://doi.org/10.1145/3397481.3450662

Tafjord, O., Mishra, B. D., & Clark, P. (2021). ProofWriter: Generating implications, proofs, and abductive statements over natural language. In C. Zong, F. Xia, W. Li, & R. Navigli (Eds.), *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021* (pp. 3621–3634). Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.findings-acl.317

Tan, R., Plummer, B., & Saenko, K. (2020). Detecting cross-modal inconsistency to defend against neural fake news. In B. Webber, T. Cohn, Y. He, & Y. Liu (Eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 2081–2106). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.emnlp-main.163

Thieme, A., Cutrell, E., Morrison, C., Taylor, A., & Sellen, A. (2020). Interpretability as a dynamic of human-AI collaboration. *ACM Interactions*, *27*(5), 40–45. https://doi.org/10.1145/3411286

Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., & Lample, G. (2023). *LLaMA: Open and efficient foundation language models*. ArXiv. https://doi.org/10.48550/arXiv.2302.13971

Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Ferrer, C. C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., . . . Scialom, T. (2023). *Llama 2: Open foundation and fine-tuned chat models*. ArXiv. https://doi.org/10.48550/arXiv.2307.09288

Turpin, M., Michael, J., Perez, E., & Bowman, S. R. (2023). *Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting*. ArXiv. https://doi.org/10.48550/arXiv.2305.04388

Ustun, B., Spangher, A., & Liu, Y. (2019). Actionable recourse in linear classification. In *FAT* '19: Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 10–19). Association for Computing Machinery. https://doi.org/10.1145/3287560.3287566

Van Der Bles, A. M., Van Der Linden, S., Freeman, A. L., Mitchell, J., Galvao, A. B., Zaval, L., & Spiegelhalter, D. J. (2019). Communicating uncertainty about facts, numbers and science. *Royal Society Open Science*, *6*(5), Article 181870. https://doi.org/10.1098/rsos.181870

Vasconcelos, H., Bansal, G., Fourney, A., Liao, Q. V., & Vaughan, J. W. (2023). *Generation probabilities are not enough: Exploring the effectiveness of uncertainty highlighting in AI-powered code completions*. ArXiv. https://doi.org/10.48550/arXiv.2302.07248

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, *30*, 5998–6008. https://papers.nips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html

Vaughan, J. W., & Wallach, H. (2021). A human-centered agenda for intelligible machine learning. In M. Pelillo & T. Scantamburlo (Eds.), *Machines we trust: Perspectives on dependable AI*. MIT Press. https://doi.org/10.7551/mitpress/12186.003.0014

Vereschak, O., Bailly, G., & Caramiaux, B. (2021). How to evaluate trust in AI-assisted decision making? A survey of empirical methodologies. *Proceedings of the ACM on Human-Computer Interaction*, *5*(CSCW2), Article 327. https://doi.org/10.1145/3476068

Vig, J., Gehrmann, S., Belinkov, Y., Qian, S., Nevo, D., Singer, Y., & Shieber, S. (2020). Investigating gender bias in language models using causal mediation analysis. *Advances in Neural Information Processing Systems*, *33*, 12388–12401. https://proceedings.neurips.cc/paper/2020/hash/92650b2e92217715fe312e6fa7b90d82-Abstract.html

Wang, D., Zhang, W., & Lim, B. Y. (2021). Show or suppress? Managing input uncertainty in machine learning model explanations. *Artificial Intelligence*, *294*, Article 103456. https://doi.org/10.1016/j.artint.2021.103456

Wang, X., & Yin, M. (2021). Are explanations helpful? A comparative study of the effects of explanations in AI-assisted decision-making. In *26th International Conference on Intelligent User Interfaces* (pp. 318–328). Association for Computing Machinery. https://doi.org/10.1145/3397481.3450650

Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., Yogatama, D., Bosma, M., Zhou, D., Metzler, D., Chi, E. H., Hashimoto, T., Vinyals, O., Liang, P., Dean, J., & Fedus W. (2022). *Emergent abilities of large language models*. ArXiv. https://doi.org/10.48550/arXiv.2206.07682

Wei, J., Wang, X., Schuurmans, D., Bosma, M., brian ichter, Xia, F., Chi, E. H., Le, Q. V., & Zhou, D. (2022). Chain of thought prompting elicits reasoning in large language models. In S. Koyejo, S.

Mohamed, A. Agarwal, D. Belgrave, K. Cho, & A. Oh (Eds.), *Advances in Neural Information Processing Systems* (Vol. 35, pp. 24824–24837). Curran Associates. https://proceedings.neurips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html

Weidinger, L., Uesato, J., Rauh, M., Griffin, C., Huang, P.-S., Mellor, J., Glaese, A., Cheng, M., Balle, B., Kasirzadeh, A., Biles, C., Brown, S., Kenton, Z., Hawkins, W., Stepleton, T., Birhane, A., Hendricks, L. A., Rimell, L., Isaac, W., . . . Gabriel, I. (2022). Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 214–229). Association for Computing Machinery. https://doi.org/10.1145/3531146.3533088

Wickens, C. D., Clegg, B. A., Vieane, A. Z., & Sebok, A. L. (2015). Complacency and automation bias in the use of imperfect automation. *Human Factors*, *57*(5), 728–739. https://doi.org/10.1177/0018720815581940

Wiegreffe, S., & Pinter, Y. (2019). Attention is not not explanation. In K. Inui, J. Jiang, V. Ng, & X. Wan (Eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (pp. 11–20). Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-1002

Wilfong, J. D. (2006). Computer anxiety and anger: The impact of computer use, computer experience, and self-efficacy beliefs. *Computers in Human Behavior*, *22*(6), 1001–1011. https://doi.org/10.1016/j.chb.2004.03.020

Wu, T., Ribeiro, M. T., Heer, J., & Weld, D. S. (2019). Errudite: Scalable, reproducible, and testable error analysis. In A. Korhonen, D. Traum, & L. Màrquez (Eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 747–763). Association for Computational Linguistics. https://doi.org/10.18653/v1/P19-1073

Wu, T., Terry, M., & Cai, C. J. (2022). AI chains: Transparent and controllable human-AI interaction by chaining large language model prompts. In S. Barbosa, C. Lamps, C. Appert, D. A. Shamma, S. Drucker J. Williamson, & K. Yatani (Eds.), *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (Article 385). Association for Computing Machinery. https://doi.org/10.1145/3491102.3517582

Yang, Q., Hao, Y., Quan, K., Yang, S., Zhao, Y., Kuleshov, V., & Wang, F. (2023). Harnessing biomedical literature to calibrate clinicians' trust in AI decision support systems. In A. Schmidt, K. Väänänen, T. Goyal, P. O. Kristensson, A. Peters, S. Mueller, J. R. Williamson, & M. L. Wilson (Eds.), *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Article 14). Association for Computing Machinery. https://doi.org/10.1145/3544548.3581393

Yi, K., Wu, J., Gan, C., Torralba, A., Kohli, P., & Tenenbaum, J. (2018). Neural-symbolic VQA: Disentangling reasoning from vision and language understanding. *Advances in Neural Information Processing Systems*, *31*, 1031–1042. https://papers.nips.cc/paper_files/paper/2018/hash/5e388103a391daabe3de1d76a6739ccd-Abstract.html

Yin, M., Vaughan, J. W., & Wallach, H. (2019). Understanding the effect of accuracy on trust in machine learning models. In *Proceedings of the 2019 CHI Conference on Human Factors in*

Computing Systems (Paper 279). Association for Computing Machinery. https://doi.org/10.1145/3290605.3300509

Yu, N., Skripniuk, V., Abdelnabi, S., & Fritz, M. (2021). Artificial fingerprinting for generative models: Rooting deepfake attribution in training data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (pp. 14448–14457). IEEE. https://doi.org/10.1109/ICCV48922.2021.01418

Zellers, R., Holtzman, A., Rashkin, H., Bisk, Y., Farhadi, A., Roesner, F., & Choi, Y. (2019). Defending against neural fake news. *Advances in Neural Information Processing Systems*, *32*, 9054–9065. https://papers.nips.cc/paper_files/paper/2019/hash/3e9f0fc9b2f89e043bc6233994dfcf76-Abstract.html

Zhang, T., Kishore, V., Wu, F., Weinberger, K. Q., & Artzi, Y. (2019). *BERTScore: Evaluating text generation with BERT*. ArXiv. https://doi.org/10.48550/arXiv.1904.09675

Zhang, Y., Liao, Q. V., & Bellamy, R. K. (2020). Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 295–305). Association for Computing Machinery. https://doi.org/10.1145/3351095.3372852

Zhou, J., Zhang, Y., Luo, Q., Parker, A. G., & Choudhury, M. D. (2023). Synthetic lies: Understanding AI-generated misinformation and evaluating algorithmic and human solutions. In A. Schmidt, K. Väänänen, T. Goyal, P. O. Kristensson, A. Peters, S. Mueller, J. R. Williamson, & M. L. Wilson (Eds.), *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Article 436). Association for Computing Machinery. https://doi.org/10.1145/3544548.3581318

©2024 Q. Vera Liao and Jennifer Wortman Vaughan. This article is licensed under a Creative Commons Attribution (CC BY 4.0) International license, except where otherwise indicated with respect to particular material included in the article.